



FAS70 and FAS90 systems

Install and maintain

NetApp

February 06, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap-systems/fas-70-90/install-overview.html> on February 06, 2026. Always check docs.netapp.com for the latest.

Table of Contents

FAS70 and FAS90 systems	1
Install and setup	1
Installation and configuration workflow - FAS70 and FAS90	1
Installation requirements - FAS70 and FAS90	1
Prepare to install - FAS70 and FAS90	3
Install the hardware - FAS70 and FAS90	5
Cable the hardware - FAS70 and FAS90	6
Power on the storage system - FAS70 and FAS90	16
Maintain	21
Overview of the maintenance procedures - FAS70 and FAS90	21
Boot media - automated recovery	23
Boot media - manual recovery	34
Chassis	56
Controller	65
Replace a DIMM - FAS70 and FAS90	79
Replace a fan - FAS70 and FAS90	85
Replace the Flash Cache module carrier or a caching module - FAS70 and FAS90	86
Replace NVRAM - FAS70 and FAS90	94
Replace the NV battery - FAS70 and FAS90	103
I/O module	108
Hot-swap a power supply - FAS70 and FAS90	126
Replace the real-time clock battery - FAS70 and FAS90	130
Replace system management module - FAS70 and FAS90	136
Key specifications	143
Key specifications for FAS70	143
Key specifications for FAS90	145

FAS70 and FAS90 systems

Install and setup

Installation and configuration workflow - FAS70 and FAS90

To install and configure your FAS70 or FAS90 system, you review the hardware requirements, prepare your site, install and cable the hardware components, power on the system, and set up your ONTAP cluster.

1

Review installation requirements

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

2

Prepare to install the FAS70 or FAS90 storage system

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

3

Install the hardware for the FAS70 or FAS90 storage system

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

4

Cable the controllers and storage shelves for the FAS70 or FAS90 storage system

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

5

Power on the FAS70 or FAS90 storage system

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

6

Set up your cluster

After you've powered on your storage system, you [set up your cluster](#).

Installation requirements - FAS70 and FAS90

Review the equipment needed and the lifting precautions for your FAS70 or FAS90 storage system and storage shelves.

Equipment needed for install

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

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

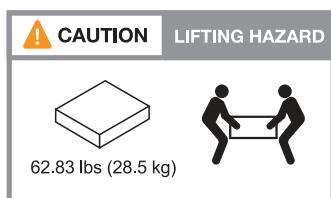
Lifting precautions

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

Storage system weight

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

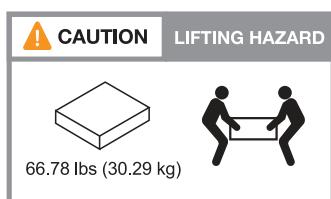
A FAS70 or FAS90 storage system can weigh up to 62.83 lbs (28.5 kg). To lift the system, use two people or a hydraulic lift.



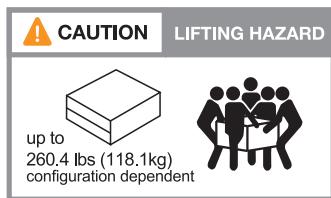
Shelf weight

Take the necessary precautions when moving or lifting your shelf.

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



A DS460C shelf can weigh up to 260.4 lbs (118.1 kg). To lift the storage shelf, you might need up to five people or a hydraulic lift. Keep all components in the storage shelf (both front and rear) to prevent unbalancing the shelf weight.



Related information

- [Safety information and regulatory notices](#)

What's next?

After you've reviewed the hardware requirements, you [prepare to install your FAS70 or FAS90 storage system](#).

Prepare to install - FAS70 and FAS90

Prepare to install your FAS70 or FAS90 storage system by getting the site ready, unpacking the boxes and comparing the contents of the boxes to the packing slip, and registering the system to access support benefits.

Step 1: Prepare the site

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

Steps

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

See the [Switch documentation](#) for installation instructions and [NetApp Hardware Universe](#) for compatibility information.

Step 2: Unpack the boxes

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

Steps

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



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

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

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

Hardware

Cables

- Bezel
- Cable management device
- Storage system
- Rail kits with instructions (optional)
- Storage shelf (if you ordered additional storage)
- Management Ethernet cables (RJ-45 cables)
- Network cables
- Power cords
- Storage cables (if you ordered additional storage)
- USB-C serial console cable

Step 3: Register your storage system

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

Steps

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

You can find the serial numbers in the following locations:

- On the packing slip
- In your confirmation email
- On each controller's System Management module

SSN: XXYYYYYYYYYY



2. Go to the [NetApp Support Site](#).

3. Determine whether you need to register your storage system:

If you are a...	Follow these steps...
Existing NetApp customer	<ol style="list-style-type: none"> Sign in with your username and password. Select Systems > My Systems. Confirm that the new serial numbers are listed. If it is not, follow the instructions for new NetApp customers.
New NetApp customer	<ol style="list-style-type: none"> Click Register Now, and create an account. Select Systems > Register Systems. Enter the storage system's serial numbers and requested details. <p>After your registration is approved, you can download any required software. The approval process might take up to 24 hours.</p>

What's next?

After you've prepared to install your FAS70 or FAS90 hardware, you [install the hardware for your FAS70 or FAS90 storage system](#).

Install the hardware - FAS70 and FAS90

After you prepare to install your FAS70 or FAS90 storage system, install the hardware for the system. First, install the rail kits. Then install and secure your platform in a cabinet or telco rack.

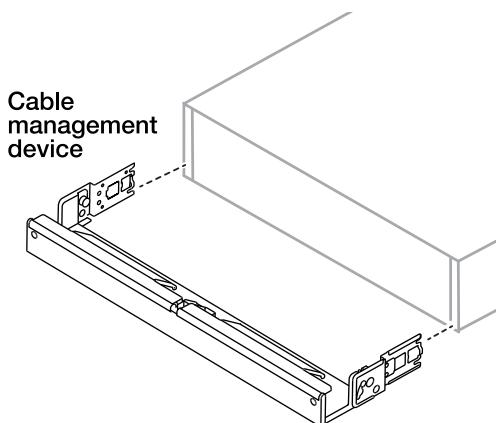
Skip this step if your cabinet is pre-populated.

Before you begin

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

Steps

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



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

If you are installing multiple shelves, place the first shelf directly above the controllers. Place the second shelf directly under the controllers. Repeat this pattern for any additional shelves.
 - b. Secure the shelf to the cabinet or telco rack using the included mounting screws.

What's next?

After you've installed the hardware for your FAS70 or FAS90 storage system, you [cable the hardware](#)

for your FAS70 or FAS90 storage system..

Cable the hardware - FAS70 and FAS90

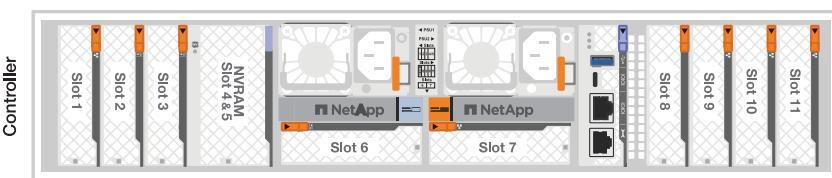
After you install the rack hardware for your FAS70 or FAS90 storage system, install the network cables for the controllers, and connect the cables between the controllers and storage shelves.

Before you begin

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

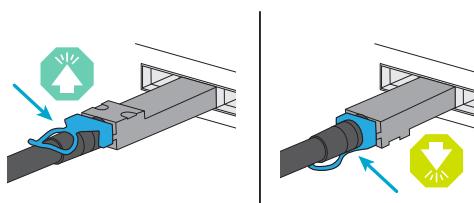
About this task

- These procedures show common configurations. The specific cabling depends on the components ordered for your storage system. For comprehensive configuration and slot priority details, see [NetApp Hardware Universe](#).
- The I/O slots on FAS70 and FAS90 controllers are numbered 1 through 11.



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

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



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

Step 1: Connect the storage controllers to your network

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



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

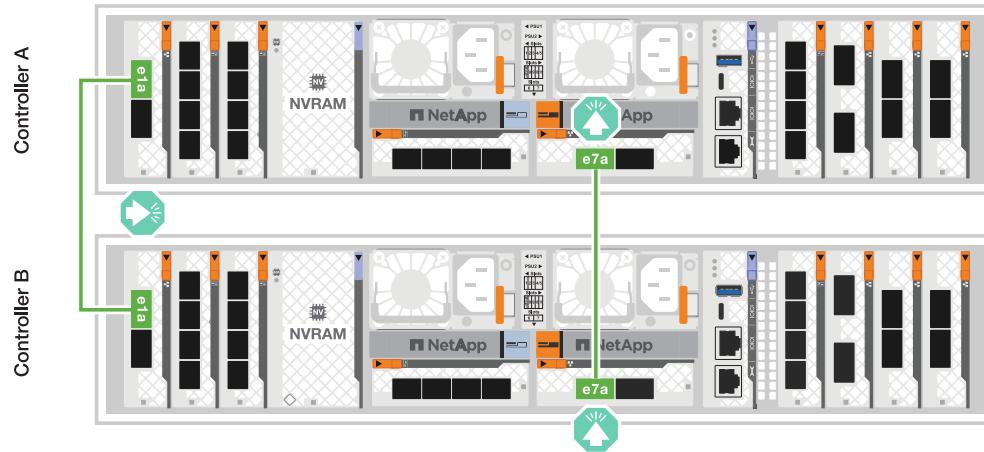
Switchless cluster cabling

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

Steps

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

Cluster/HA interconnect cables



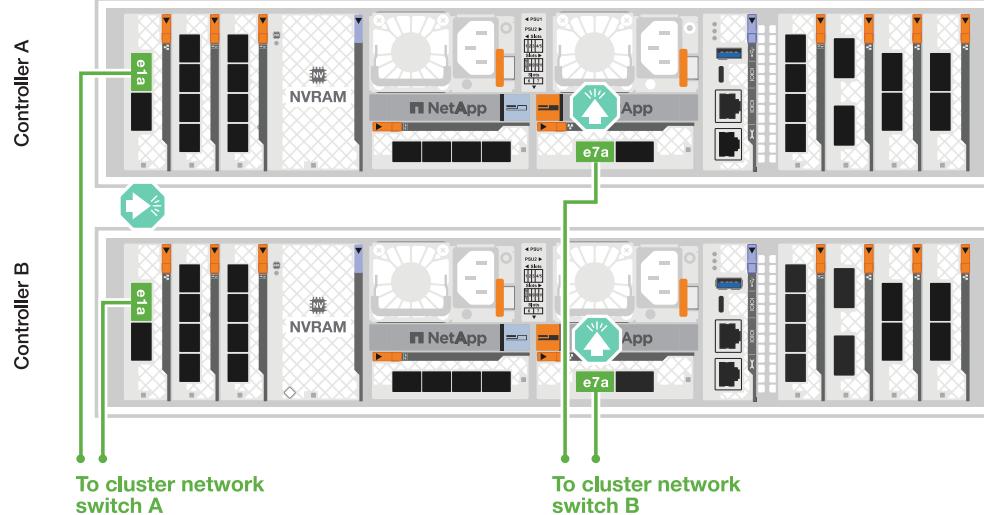
Switched cluster cabling

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

Steps

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

100 GbE cable



Step 2: Cable the host network connections

Connect the Ethernet module ports to your host network.

The following are some typical host network cabling examples. See [NetApp Hardware Universe](#) for your specific system configuration.

Steps

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

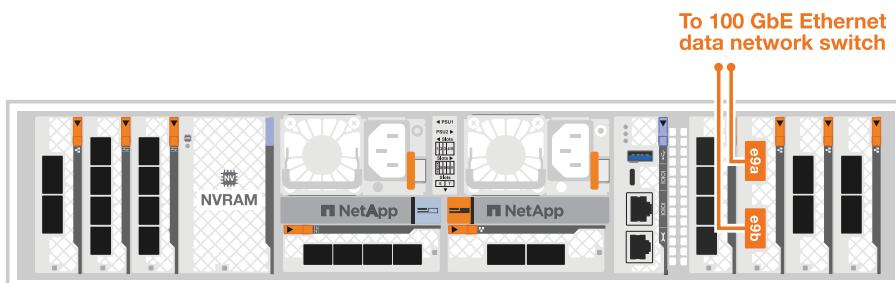


Do not use ports e1b and e7b ports for host network connections. Use a separate host card.

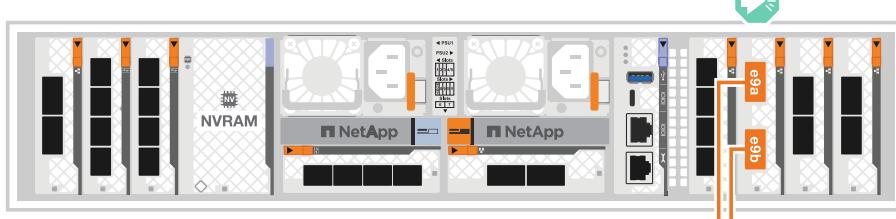
100 GbE cable



Controller A



Controller B

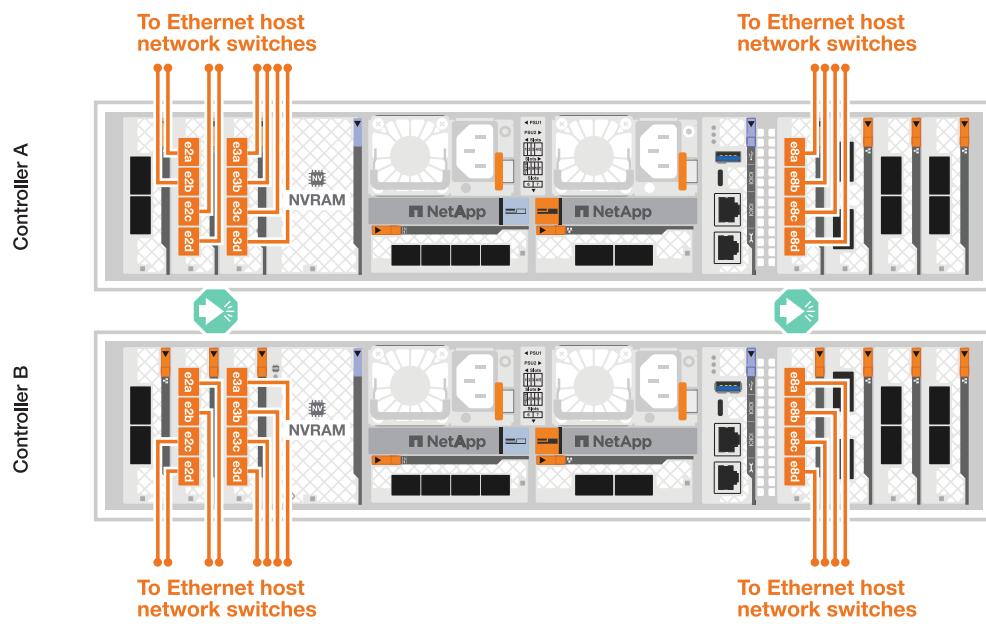


To 100 GbE Ethernet data network switch

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

4-ports, 10/25 GbE Host



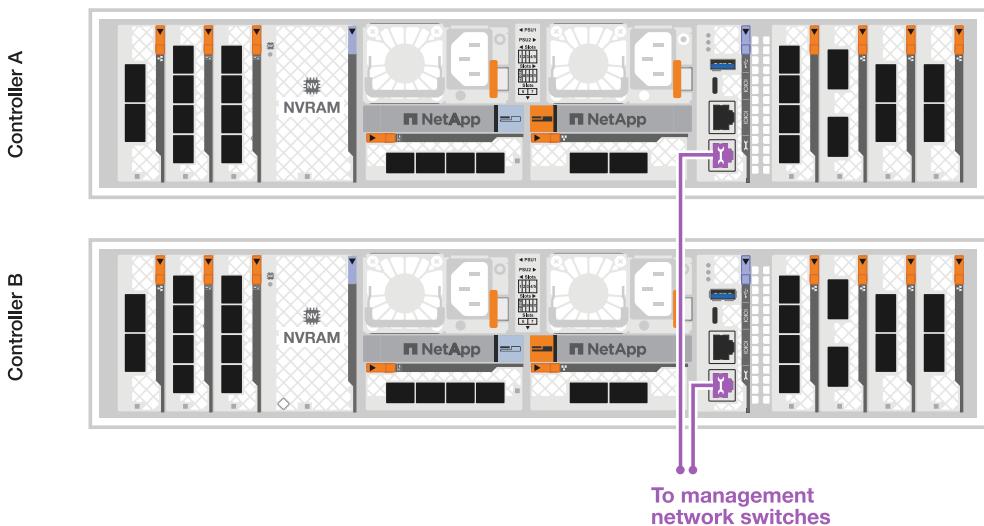


Step 3: Cable the management network connections

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



1000BASE-T RJ-45 cables



Do not plug in the power cords yet.

Step 4: Cable the shelf connections

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

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

For additional SAS shelf cabling guidance, see [SAS cabling rules and concepts - shelves with IOM12/IOM12B modules](#).

About this task

The FAS70 and FAS90 storage systems supports DS212C, DS224C, DS460C, and NS224 shelves with either the NSM100 or NSM100B module.

The major differences between the NS224 modules are:

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

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

Option 1: One NS224 storage shelf

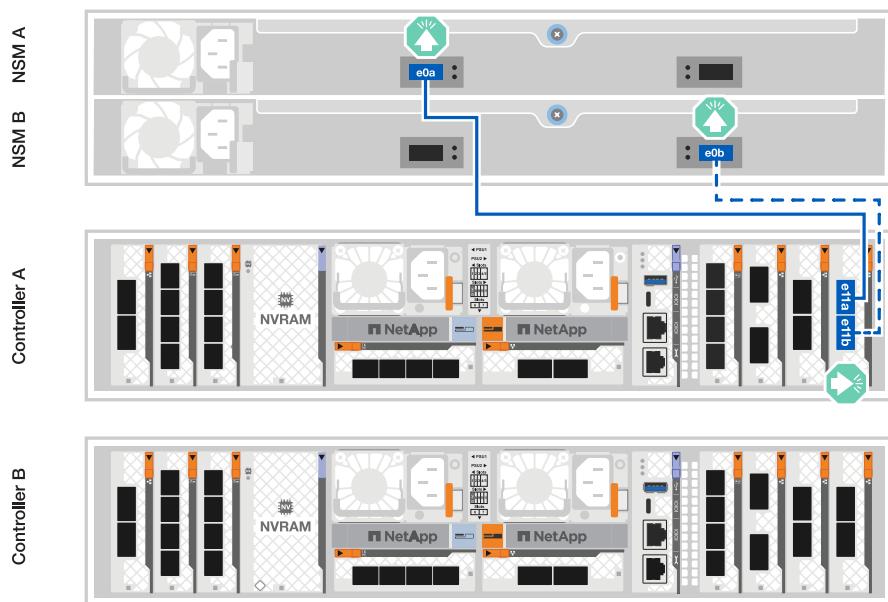
Connect each controller to the NSM modules on the NS224 shelf. The graphics show controller A cabling in blue and controller B cabling in yellow.

100 GbE QSFP28 copper cables

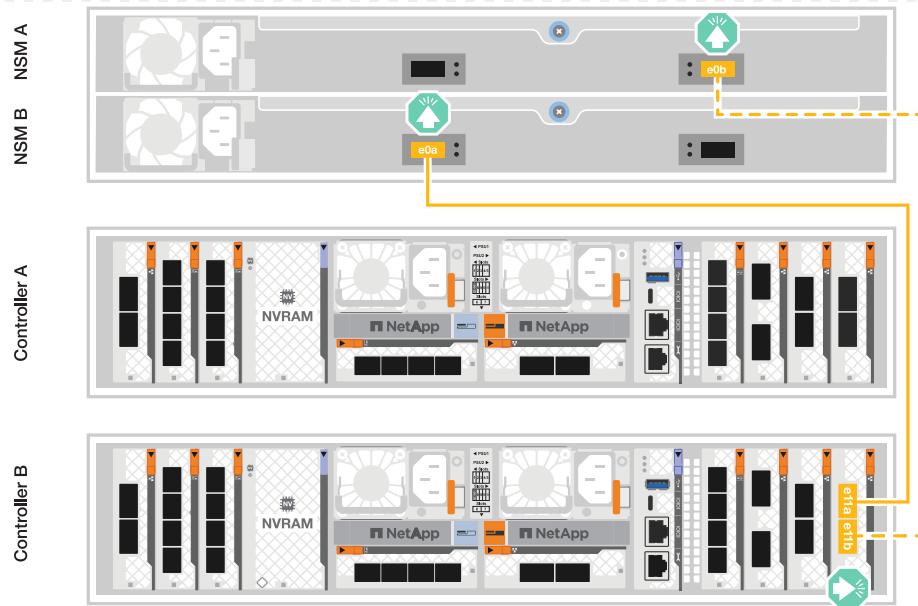


Steps

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



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



Option 2: Two NS224 storage shelves

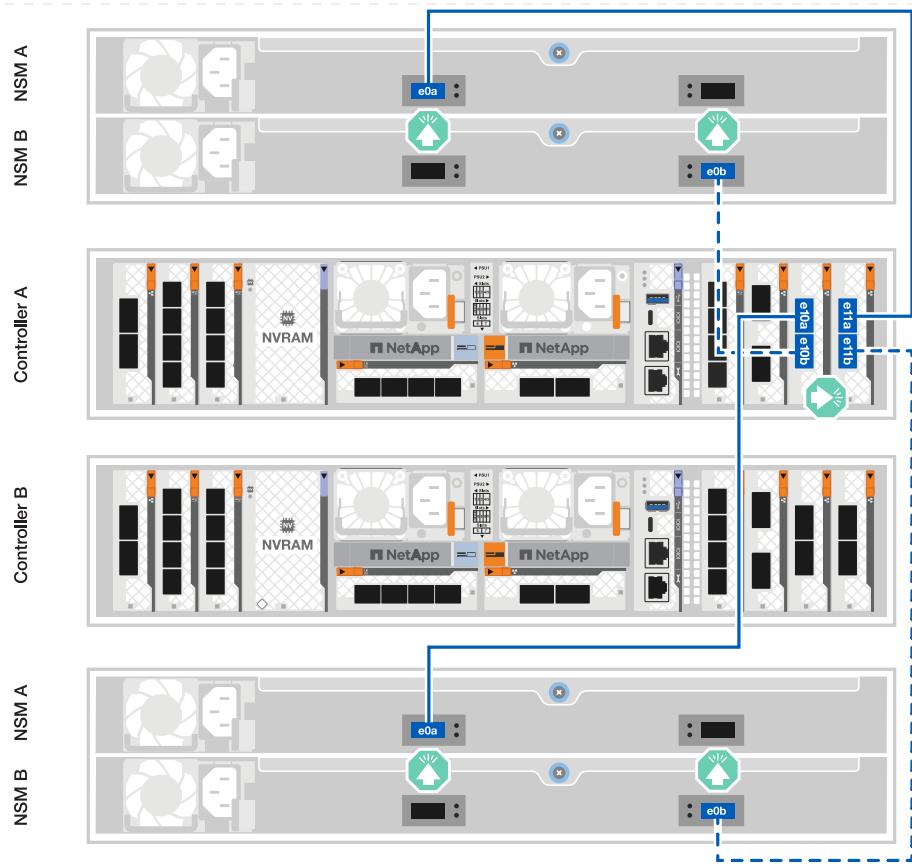
Cable each controller to the NSM modules on both NS224 shelves. The graphics show controller A cabling in blue and controller B cabling in yellow.

100 GbE QSFP28 copper cables



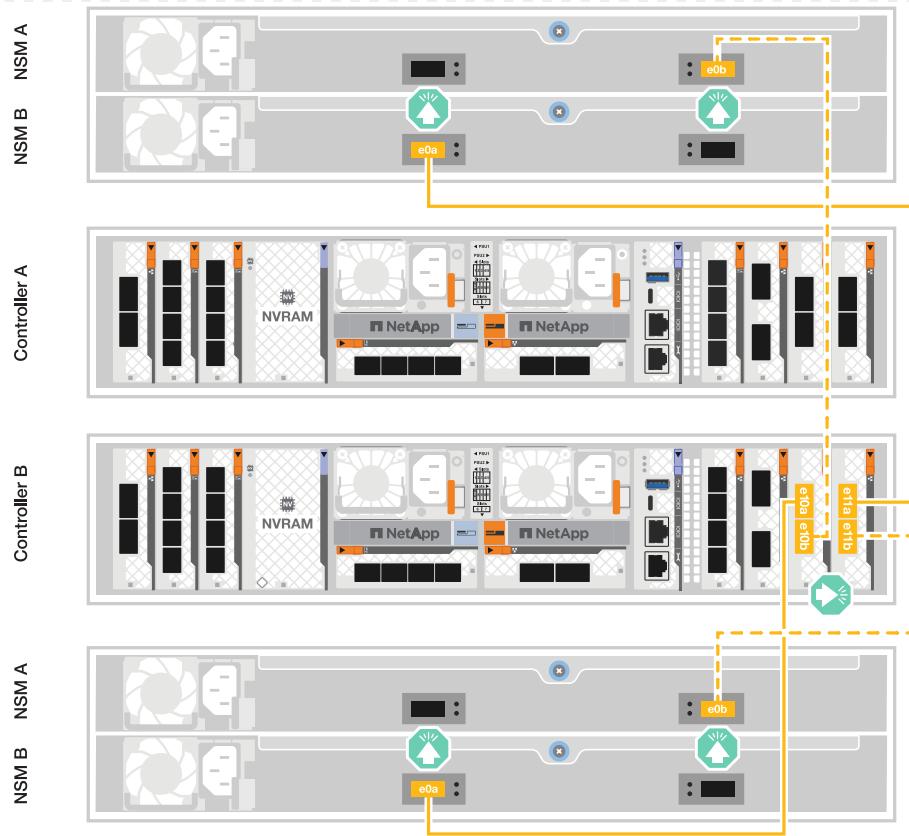
Steps

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



2. On controller B, connect the following ports:

- Connect port e11a to shelf 1 NSM B port e0a.
- Connect port e11b to shelf 2 NSM A port e0b.
- Connect port e10a to shelf 2 NSM B port e0a.
- Connect port e10b to shelf 1 NSM A port e0b.



Option 3: Two DS460C shelves

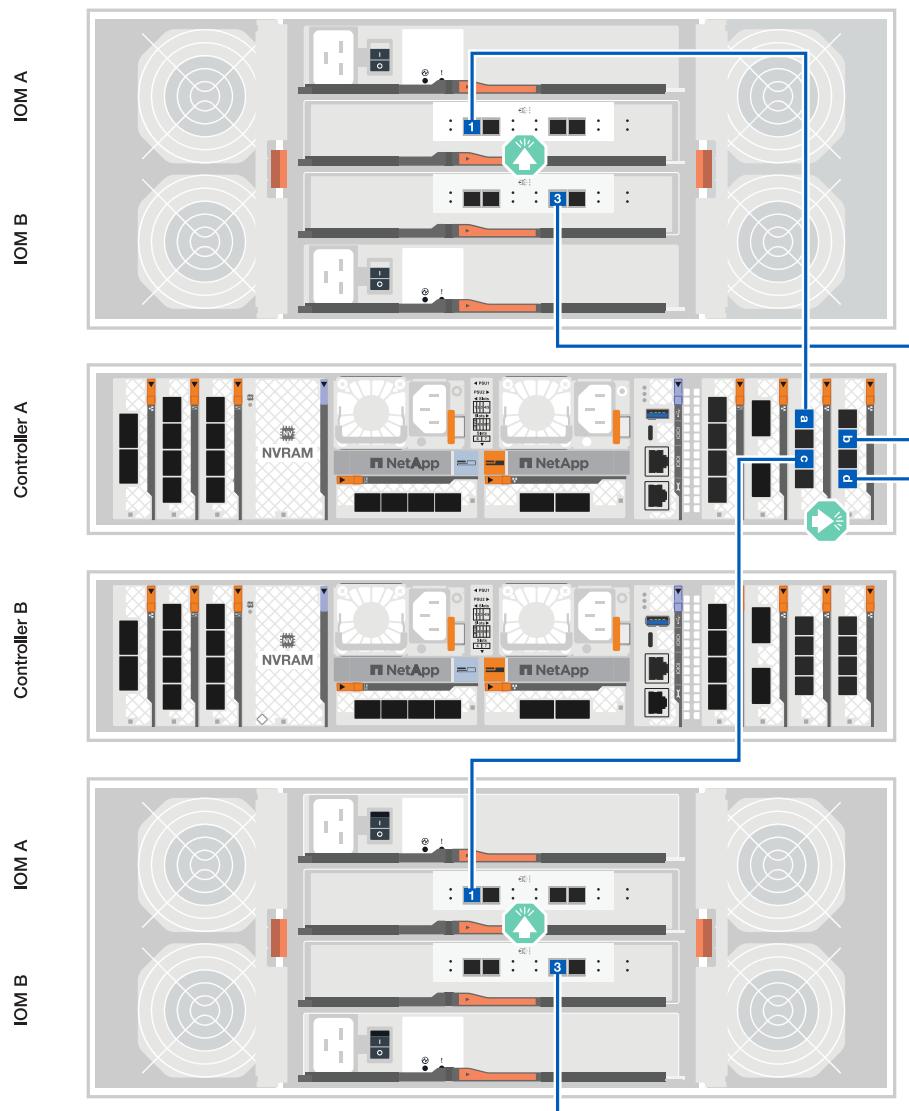
Cable each controller to the IOM modules on both DS460C shelves. The graphics show controller A cabling in blue and controller B cabling in yellow.

mini-SAS HD cable



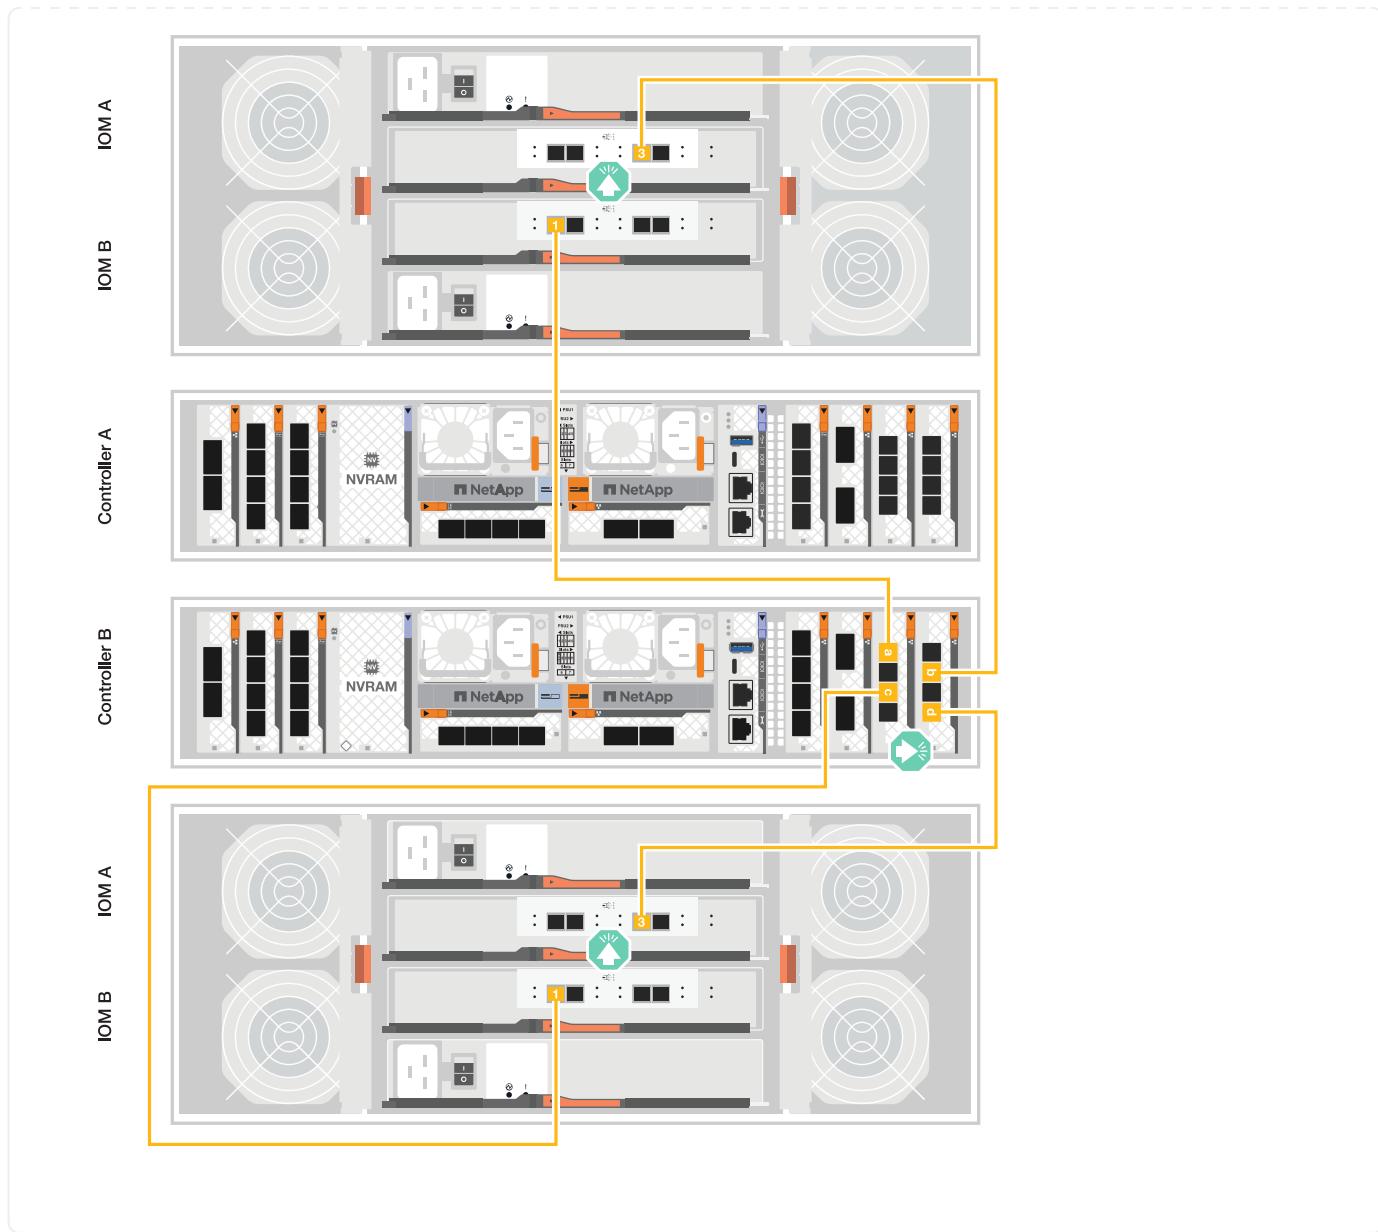
Steps

1. On controller A, cable the following connections:
 - a. Connect port e10a to shelf 1 IOM A port 1.
 - b. Connect port e10c to shelf 2 IOM A port 1
 - c. Connect port e11b to shelf 1 IOM B port 3.
 - d. Connect port e11d to shelf 2 IOM B port 3.



2. On controller B, cable the following connections:

- Connect port e10a to shelf 1 IOM B port 1.
- Connect port e10c to shelf 2 IOM B port 1.
- Connect port e11b to shelf 1 IOM A port 3.
- Connect port e11d to shelf 2 IOM A port 3.



What's next?

After you've cabled the hardware for your FAS70 or FAS90 system, you [power on the FAS70 or FAS90 storage system](#).

Power on the storage system - FAS70 and FAS90

After you install the rack hardware for your FAS70 or FAS90 storage system and install the cables for the controllers and storage shelves, you should power on your storage shelves and controllers.

Step 1: Power on the shelf and assign shelf ID

Option 1: NS224 shelves

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

Before you begin

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

About this task

- A valid shelf ID is 01 through 99.

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

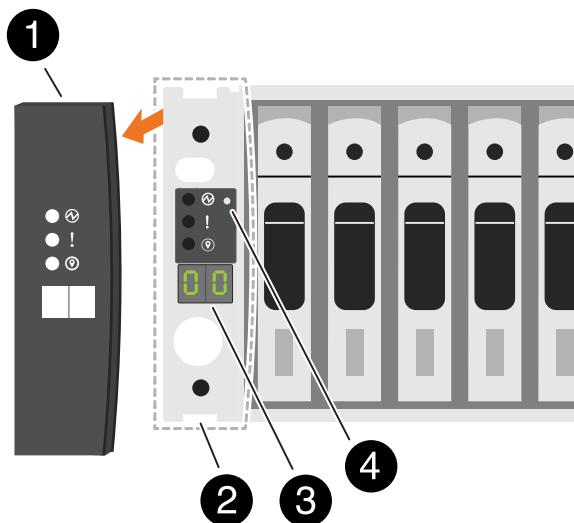
- You must power cycle a shelf (unplug both power cords, wait the appropriate amount of time, and then plug them back in) for the shelf ID to take effect.

Steps

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

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

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



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

3. Change the first number of the shelf ID:

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

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



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

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

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

The first number continues to blink.

4. Change the second number of the shelf ID:

- a. Press and hold the button until the second number on the digital display blinks.

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

The first number on the digital display stops blinking.

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

The second number continues to blink.

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

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

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

6. Power-cycle the shelf for at least 10 seconds to make the shelf ID take effect.

- a. Unplug the power cord from both power supplies on the shelf.
- b. Wait 10 seconds.
- c. Plug the power cords back into the shelf power supplies to complete the power cycle.

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

7. Replace the left end cap.

Option 2: SAS shelves

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

About this task

- A valid shelf ID is 01 through 99.

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

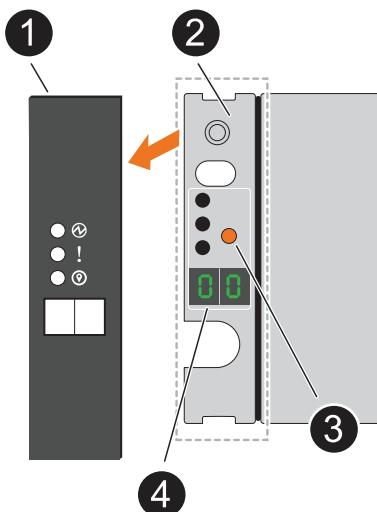
- You must power cycle a shelf (turn off the power switch on each of the power supplies of the SAS shelf, wait the appropriate amount of time, and then switch the power back on) for the shelf ID to take effect.

Steps

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

The shelf powers on and boots automatically when powered on.

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



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

3. Change the first number of the shelf ID:

- a. Press and hold the shelf ID button until the first number on the digital display blinks, and then release the button.

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



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

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

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

The first number continues to blink.

4. Change the second number of the shelf ID:

- a. Press and hold the button until the second number on the digital display blinks.

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

The first number on the digital display stops blinking.

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

The second number continues to blink.

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

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

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

6. Power-cycle the shelf for at least 10 seconds to make the shelf ID take effect.

- a. Turn off the power switch on each of the power supplies.

- b. Wait 10 seconds.

- c. Turn on the power switch on each of the power supplies to complete the power cycle.

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

7. Replace the left end cap.

Step 2: Power on the controllers

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

Steps

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

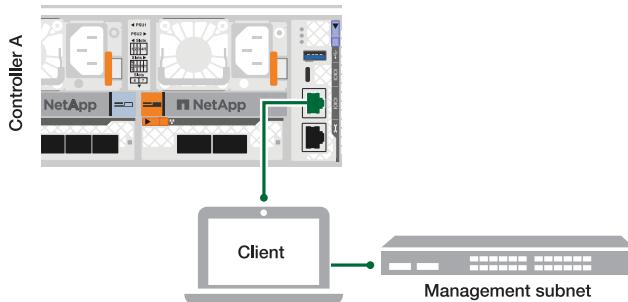


See your laptop's online help for instructions on how to configure the serial console port.

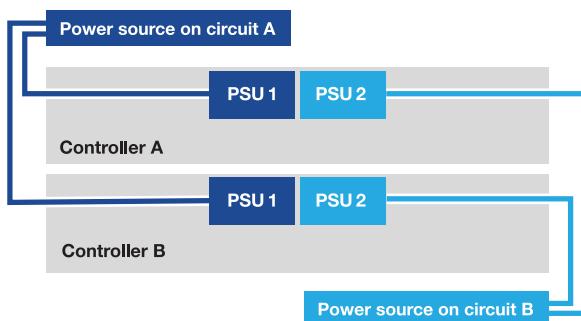
- b. Using the console cable provided with your storage system, connect one end of the console cable to

your laptop and the other end to the serial console port on controller A.

- c. Connect the laptop to the switch on the management subnet.



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



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

What's next?

After you've turned on your FAS70 or FAS90 storage system, you [set up your cluster](#).

Maintain

Overview of the maintenance procedures - FAS70 and FAS90

Maintain the hardware of your FAS70 or FAS90 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the FAS70 or FAS90 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the FAS70 and FAS90 storage system, you can perform maintenance procedures on the following components.

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

Boot media - automated recovery

Boot media automated recovery workflow - FAS70 and FAS90

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your FAS70 or FAS90 storage system.

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

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

1

[Review the boot media requirements](#)

Review the requirements for boot media replacement.

2

[Shut down the controller](#)

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

3

[Replace the boot media](#)

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

4

[Restore the image on the boot media](#)

Restore the ONTAP image from the partner controller.

5

[Return the failed part to NetApp](#)

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

Requirements for automated boot media recovery - FAS70 and FAS90

Before replacing the boot media in your FAS70 or FAS90 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming the cluster ports on the impaired controller are working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

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

Review the following requirements.

- You must replace the failed boot media with a replacement boot media you received from NetApp.
- The cluster ports are used to communicate between the two controllers during the automated boot recovery process. Make sure that the cluster ports on the impaired controller are working properly.
- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
 - /cfcards/kmip/servers.cfg
 - /cfcards/kmip/certs/client.crt
 - /cfcards/kmip/certs/client.key
 - /cfcards/kmip/certs/CA.pem
- Understand the controller terminology used in this procedure:
 - The *impaired controller* is the controller on which you are performing maintenance.
 - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you [shut down the controller](#).

Shut down the controller for automated boot media recovery - FAS70 and FAS90

Shut down the impaired controller in your FAS70 or FAS90 storage system to prevent data loss and maintain system stability during the automatic boot media recovery process.

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

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

- Enter the following command from the console of the healthy controller:

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

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

3. Take the impaired controller to the LOADER prompt:

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

What's next

After you shut down the impaired controller, you [replace the boot media](#).

Replace the boot media for automated boot recovery - FAS70 and FAS90

The boot media in your FAS70 or FAS90 storage system stores essential firmware and configuration data. The replacement process involves removing the System Management module, removing the impaired boot media, installing the replacement boot media in the System Management module, and then reinstalling the System Management module.

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

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

Steps

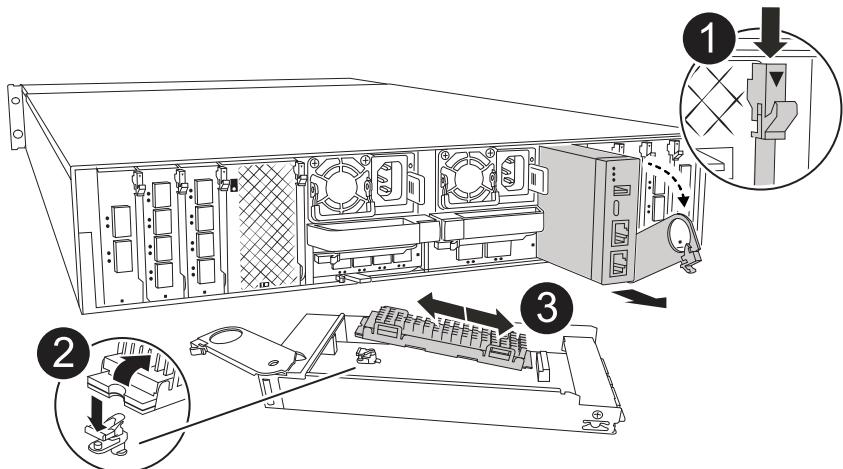
- If you are not already grounded, properly ground yourself.
- Unplug the power supply cables from the PSUs.



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

- Remove the System Management module:

- a. Remove any cables connected to the System Management module. Make sure that you label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
 - b. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
 - c. Depress the System Management cam button.
 - d. Rotate the cam latch down as far as it will go.
 - e. Remove the System Management module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
 - f. Place the System Management module on an anti-static mat, so that the boot media is accessible.
4. Remove the boot media from the management module:



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

- a. Press the blue locking button.
 - b. Rotate the boot media up, slide it out of the socket, and set it aside.
5. Install the replacement boot media into the System Management module:
- a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
 - b. Rotate the boot media down toward the locking button.
 - c. Push the locking button, rotate the boot media all the way down and then release the locking button.
6. Reinstall the System Management module:
- a. Align the module with the edges of the enclosure slot opening.
 - b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.

7. Rotate the cable management tray up to the closed position.
 - a. Recable the System Management module.
8. Plug the power cables into the power supplies and reinstall the power cable retainer.

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

What's next

After physically replacing the impaired boot media, [restore the ONTAP image from the partner node](#).

Automated boot media recovery from the partner node - FAS70 and FAS90

After installing the new boot media device in your FAS70 or FAS90 storage system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

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

Before you begin

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

Steps

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

```
boot_recovery -partner
```

The screen displays the following message:

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

2. Monitor the boot media install recovery process.

The process completes and displays the Installation complete message.

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

If you see this message...	Do this...
key manager is not configured. Exiting.	<p>Encryption is not installed on the system.</p> <ol style="list-style-type: none"> Wait for the login prompt to display. Log into the node and give back the storage: <pre>storage failover giveback -ofnode <i>impaired_node_name</i></pre> <ol style="list-style-type: none"> Go to re-enabling automatic giveback if it was disabled.
key manager is configured.	Encryption is installed. Go to restoring the key manager .



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

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

Onboard Key Manager (OKM)

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

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

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

- a. Enter **y** at the prompt to confirm you want to start the OKM recovery process.
- b. Enter the passphrase for onboard key management when prompted.
- c. Enter the passphrase again when prompted to confirm.
- d. Enter the backup data for onboard key manager when prompted.

Show example of passphrase and backup data prompts

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

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

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

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

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

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

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

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

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

Go to [re-enabling automatic giveback](#) if it was disabled.

External Key Manager (EKM)

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

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

- a. Enter the EKM configuration settings when prompted:

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

Show example of client certificate contents

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

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

Show example of client key file contents

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

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

Show example of KMIP server file contents

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

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

Show example of server configuration file contents

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

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

Show example of ONTAP Cluster UUID prompt

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

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

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

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

Show example of temporary network setting prompts

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

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

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

b. Verify the key restoration status:

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

Show example of key recovery error and warning messages

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

Uptime: 11m32s
System halting...

LOADER-B>
```

- c. After the node reboots, verify that the system is back online and operational.
- d. Return the controller to normal operation by giving back its storage:

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

Go to [re-enabling automatic giveback](#) if it was disabled.

5. If automatic giveback was disabled, reenable it:

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

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

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

What's next

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

Return the failed boot media part to NetApp - FAS70 and FAS90

If a component in your FAS70 or FAS90 storage system fails, return the failed part to NetApp. See the [Part Return and Replacements](#) page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - FAS70 and FAS90

The manual recovery of the boot image involves using a USB drive to reinstall ONTAP onto the FAS70 or FAS90 system's replacement boot media. You must download the appropriate ONTAP recovery image from the NetApp Support Site and copy it to a USB drive. This prepared USB drive is then used to perform the recovery and restore the system to operational status.

If your system is running in ONTAP 9.17.1 and later, use the [automatic boot recovery procedure](#).

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

1

[Review boot media replacement requirements](#)

Review the requirements for replacing the boot media.

2

[Check onboard encryption keys](#)

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

3

[Shut down the impaired controller](#)

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

4

[Replace the boot media](#)

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

5

[Boot the recovery image](#)

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

6

[Restore encryption](#)

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

7

[Return the failed part to NetApp](#)

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

Requirements for manual boot media recovery - FAS70 and FAS90

Before replacing the boot media in your FAS70 or FAS90 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

USB flash drive

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

File preparation

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

Component replacement

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

Controller identification

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

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

What's next?

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

Check encryption support for manual boot media recovery - FAS70 and FAS90

To ensure data security on your FAS70 or FAS90 storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

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

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

Steps

1. Check if your ONTAP version supports encryption:

```
version -v
```

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

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

- If NVE is supported: Download the ONTAP image with NetApp Volume Encryption

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



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

Step 2: Verify key manager status and back up configuration

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

Steps

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

ONTAP version	Run this command
ONTAP 9.14.1 or later	<p><code>security key-manager keystore show</code></p> <ul style="list-style-type: none"> • If EKM is enabled, EKM is listed in the command output. • If OKM is enabled, OKM is listed in the command output. • If no key manager is enabled, No key manager keystores configured is listed in the command output.
ONTAP 9.13.1 or earlier	<p><code>security key-manager show-key-store</code></p> <ul style="list-style-type: none"> • If EKM is enabled, external is listed in the command output. • If OKM is enabled, onboard is listed in the command output. • If no key manager is enabled, No key managers configured is listed in the command output.

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

If no key manager is configured:

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

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

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

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

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

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

External Key Manager (EKM)

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

If all keys show true in the Restored column:

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

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

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

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

If the command fails, contact NetApp Support.

- Verify that all authentication keys are restored:

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

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

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

Onboard Key Manager (OKM)

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

If all keys show true in the Restored column:

- Back up the OKM information:

- Switch to advanced privilege mode:

```
set -priv advanced
```

Enter y when prompted to continue.

- Display the key management backup information:

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

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

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

- Return to admin mode:

```
set -priv admin
```

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

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

a. Synchronize the onboard key manager:

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

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



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

b. Verify all authentication keys are restored:

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

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

c. Back up the OKM information:

i. Switch to advanced privilege mode:

```
set -priv advanced
```

Enter `y` when prompted to continue.

ii. Display the key management backup information:

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

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

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

iv. Return to admin mode:

```
set -priv admin
```

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

What's next?

After checking the encryption key support and status on the boot media, you need to [shut down the controller](#).

Shut down the controller for manual boot media recovery - FAS70 and FAS90

Shut down the impaired controller in your FAS70 or FAS90 storage system to prevent data loss and maintain system stability during the manual boot media recovery process.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

3. Take the impaired controller to the LOADER prompt:

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

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: Controller is in a MetroCluster

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

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

```
metrocluster node show
```

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

If the impaired controller is displaying...	Then...
System prompt or password prompt (enter system password)	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

What's next?

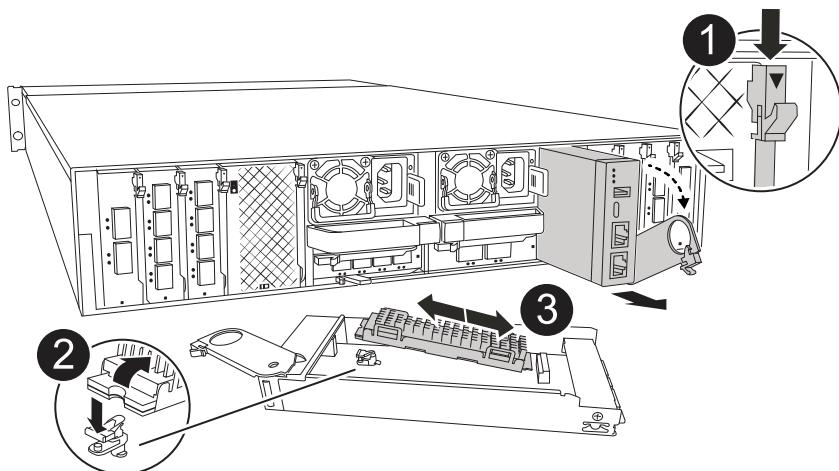
After shutting down the controller, you need to [replace the boot media](#).

Replace the boot media and prepare for manual boot recovery - FAS70 and FAS90

The boot media in your FAS70 or FAS90 system stores essential firmware and configuration data. The replacement process involves removing the System Management module, removing the impaired boot media, installing the replacement boot media, and then manually transferring the ONTAP image to the replacement boot media using a USB flash drive.

Step 1: Replace the boot media

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



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

Steps

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



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

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

Step 2: Transfer the ONTAP image to the boot media

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

Before you begin

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

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

- If NVE is supported by your version of ONTAP, download the image with NetApp Volume Encryption, as indicated in the download button.
- If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection between the node management ports of the controllers (typically the e0M interfaces).

Steps

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

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

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

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

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

3. Plug the power cables back into the PSU.

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

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

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

What's next?

After replacing the boot media, you need to [boot the recovery image](#).

Manual boot media recovery from a USB drive - FAS70 and FAS90

After installing the new boot media device in your FAS70 or FAS90 system, you can boot the recovery image manually from a USB drive to restore the configuration from the partner node.

Before you begin

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

Steps

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

```
boot_recovery
```

The recovery image is downloaded from the USB flash drive.

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

ONTAP 9.16.0 or earlier

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

- a. **On the impaired controller:** Press Y when you see Do you want to restore the backup configuration now?
- b. **On the impaired controller:** If prompted, press Y to overwrite `/etc/ssh/ssh_host_ecdsa_key`.
- c. **On the partner controller:** Set the impaired controller to advanced privilege level:

```
set -privilege advanced
```

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

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



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

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

```
set -privilege admin
```

- f. **On the impaired controller:** Press Y when you see Was the restore backup procedure successful?
- g. **On the impaired controller:** Press Y when you see ...would you like to use this restored copy now?
- h. **On the impaired controller:** Press Y when prompted to reboot, then press `Ctrl-C` when you see the Boot Menu.
- i. **On the impaired controller:** Do one of the following:
 - If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.
 - If the system uses encryption, go to [Restore encryption](#).

ONTAP 9.16.1 or later

Complete the following steps on the impaired controller:

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

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

- b. Press Y when prompted to confirm that the restore backup was successful.
- c. Press Y when prompted to use the restored configuration.
- d. Press Y when prompted to reboot the node.
- e. Press Y when prompted to reboot again, then press `Ctrl-C` when you see the Boot Menu.
- f. Do one of the following:
 - If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.

- If the system uses encryption, go to [Restore encryption](#).

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

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

6. If you disabled automatic giveback, reenable it:

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

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

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

What's next?

After booting the recovery image, you need to [restore encryption on the boot media](#).

Restore encryption keys after manual boot recovery - FAS70 and FAS90

Restore encryption on the replacement boot media in your FAS70 or FAS90 system to ensure continued data protection. The replacement process involves verifying key availability, reapplying encryption settings, and confirming secure access to your data.

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

Onboard Key Manager (OKM)

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

Before you begin

Ensure you have the following information available:

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

Steps

On the impaired controller:

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

ONTAP version	Select this option
ONTAP 9.8 or later	<p>Select option 10.</p> <p>Show example boot menu</p> <div style="border: 1px solid #ccc; padding: 10px; width: fit-content; margin: 10px auto; border-radius: 10px;"><p>Please choose one of the following:</p><ul style="list-style-type: none">(1) Normal Boot.(2) Boot without /etc/rc.(3) Change password.(4) Clean configuration and initialize all disks.(5) Maintenance mode boot.(6) Update flash from backup config.(7) Install new software first.(8) Reboot node.(9) Configure Advanced Drive Partitioning.(10) Set Onboard Key Manager recovery secrets.(11) Configure node for external key management.<p>Selection (1-11)? 10</p></div>

ONTAP version	Select this option
ONTAP 9.7 and earlier	<p>Select the hidden option <code>recover_onboard_keymanager</code></p> <p>Show example boot menu</p> <div style="border: 1px solid #ccc; padding: 10px; border-radius: 10px; background-color: #f9f9f9;"> <p style="margin: 0;">Please choose one of the following:</p> <p style="margin: 0;">(1) Normal Boot.</p> <p style="margin: 0;">(2) Boot without /etc/rc.</p> <p style="margin: 0;">(3) Change password.</p> <p style="margin: 0;">(4) Clean configuration and initialize all disks.</p> <p style="margin: 0;">(5) Maintenance mode boot.</p> <p style="margin: 0;">(6) Update flash from backup config.</p> <p style="margin: 0;">(7) Install new software first.</p> <p style="margin: 0;">(8) Reboot node.</p> <p style="margin: 0;">(9) Configure Advanced Drive Partitioning.</p> <p style="margin: 0;">Selection (1-19)?</p> <p style="margin: 0;">recover_onboard_keymanager</p> </div>

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

Show example prompt

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

4. Enter the cluster-wide passphrase twice.

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

Show example prompt

Enter the passphrase for onboard key management:

Enter the passphrase again to confirm:

5. Enter the backup information:

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

Show example prompt

Enter the backup data:

```
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA  
AA  
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA  
AA  
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA  
AA  
0123456789012345678901234567890123456789012345678901  
23  
1234567890123456789012345678901234567890123456789012  
34  
2345678901234567890123456789012345678901234567890123  
45  
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAA  
AA  
AAAAAAAAAAAAAAAAAAAAAA  
AA  
AAAAAAAAAAAAAAAAAAAAAA  
AA  
-----END  
BACKUP-----
```

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

The recovery process completes and displays the following message:

Successfully recovered keymanager secrets.

Show example prompt

```
Trying to recover keymanager secrets....  
Setting recovery material for the onboard key manager  
Recovery secrets set successfully  
Trying to delete any existing km_onboard.wkeydb file.  
  
Successfully recovered keymanager secrets.  
  
*****  
*****  
* Select option "(1) Normal Boot." to complete recovery  
process.  
*  
* Run the "security key-manager onboard sync" command to  
synchronize the key database after the node reboots.  
*****  
*****
```



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

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

Show example prompt

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

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

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

On the partner controller:

8. Giveback the impaired controller:

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

On the impaired controller:

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

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

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

Show example prompt

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

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



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

11. Verify that all keys are synced:

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

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

On the partner controller:

12. Giveback the impaired controller:

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

13. Restore automatic giveback if you disabled it:

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

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

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

External Key Manager (EKM)

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

Before you begin

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

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

Steps

On the impaired controller:

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

Show example boot menu

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

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

Show example prompt

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

4. Enter the client and server information when prompted:

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

Show example

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

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

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

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

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

The recovery process completes and displays the following message:

```
Successfully recovered keymanager secrets.
```

Show example

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

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

Show example prompt

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

6. Restore automatic giveback if you disabled it:

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

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

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

What's next?

After restoring encryption on the boot media, you need to [return the failed part to NetApp](#).

Return the failed part to NetApp - FAS70 and FAS90

If a component in your FAS70 or FAS90 system fails, return the failed part to NetApp. See the [Part Return and Replacements](#) page for further information.

Chassis

Chassis replacement workflow - FAS70 and FAS90

Get started with replacing the chassis of your FAS70 or FAS90 storage system by reviewing the replacement requirements, shutting down the controller, replacing the

chassis, and verifying system operations.

1

Review the chassis replace requirements

Review the chassis replacement requirements.

2

Prepare for chassis replace

Prepare to replace the chassis by locating the system, gathering system credentials and necessary tools, verifying the replacement chassis was received, and labeling the system cables.

3

Shut down the controller

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

4

Replace the chassis

Replace the chassis by moving the components from the impaired chassis to the replacement chassis.

5

Complete the chassis replacement

Complete the chassis replacement by bringing the controller up, giving back the controller, and returning the failed chassis to NetApp.

Requirements to replace the chassis - FAS70 and FAS90

Before replacing the chassis in your FAS70 or FAS90 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have local administrator credentials for ONTAP, the correct replacement chassis, and the necessary tools.

The chassis is the physical chassis housing all the system components such as the fans, controller/CPU unit, NVRAM12, system management module, I/O cards and blanking modules, and PSUs.

Review the following requirements.

- Make sure all other components in the system are functioning properly; if not, contact [NetApp support](#) for assistance.
- Obtain local administrator credentials for ONTAP if you don't have them.
- Make sure that you have the necessary tools and equipment for the replacement.
- You can use the chassis replacement procedure with all versions of ONTAP supported by your system.
- The chassis replacement procedure is written with the assumption that you are moving the bezel, fans, controller module, NVRAM12, system management module, I/O cards and blanking modules, and PSUs to the new chassis, and that the replacement chassis is a new component from NetApp.

What's next?

After you've reviewed the requirements to replace the chassis, you need to [prepare to replace the chassis](#).

Prepare to replace the chassis - FAS70 and FAS90

Prepare to replace the impaired chassis in your FAS70 or FAS90 system by identifying the impaired chassis, verifying the replacement components, and labeling the cables and controller module.

Step 1: Locate and monitor your system

You should open a console session and save sessions logs for future reference, and also turn on the system location LED to find the impaired chassis.

Steps

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

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

The Location LED remains lit for 30 minutes.

Step 2: Verify replacement components

You should verify that you received the necessary components, remove them from packaging, and save the packaging.

Steps

1. Before opening the packaging, you should look at the packaging label and verify:
 - Component part number.
 - Part description.
 - Quantity in the box.
2. Remove the contents from the packaging and use the packaging to returning the failed component to NetApp.

Step 3: Label the cables

You should label the cables before removing them from the I/O modules on the back of the system.

Steps

1. Label all the cables associated with the storage system. This aids recabling later in this procedure.
2. If you are not already properly grounded, ground yourself.

What's next?

After you've prepared to replace your FAS70 or FAS90 chassis hardware, you need to [shut down the controller](#).

Shut down the controller to replace the chassis - FAS70 and FAS90

Shut down the controller in your FAS70 or FAS90 storage system to prevent data loss and ensure system stability when replacing the chassis.

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

What's next?

After you've shut down the controller, you need to [replace the chassis](#).

Replace the chassis - FAS70 and FAS90

Replace the chassis of your FAS70 or FAS90 system when a hardware failure requires it. The replacement process involves removing the controller, I/O cards, NVRAM12 module, system management module, and power supply units (PSUs), installing the replacement chassis, and reinstalling the chassis components.

Step 1: Remove the PSUs and cables

You need to remove the two power supply units (PSUs) before removing the controller.

Steps

1. Remove the PSUs:

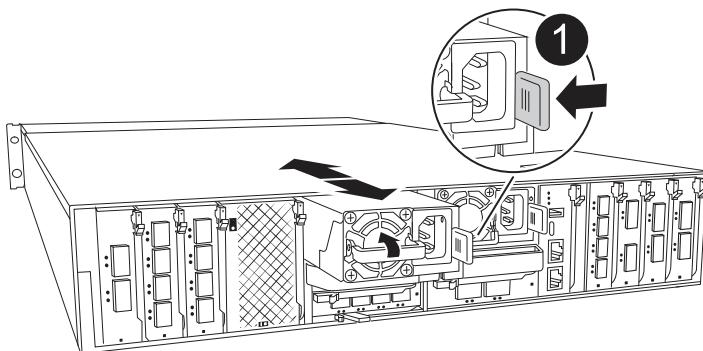
- If you are not already grounded, properly ground yourself.
- Unplug power cords from the PSUs.

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

- Remove the two PSUs from the rear of the chassis by rotating the PSU handle up so that you can pull the PSU out, press the PSU locking tab, and then pull PSU out of the chassis.



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



Terracotta PSU locking tab

- Repeat these steps for the second PSU.

2. Remove the cables:

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

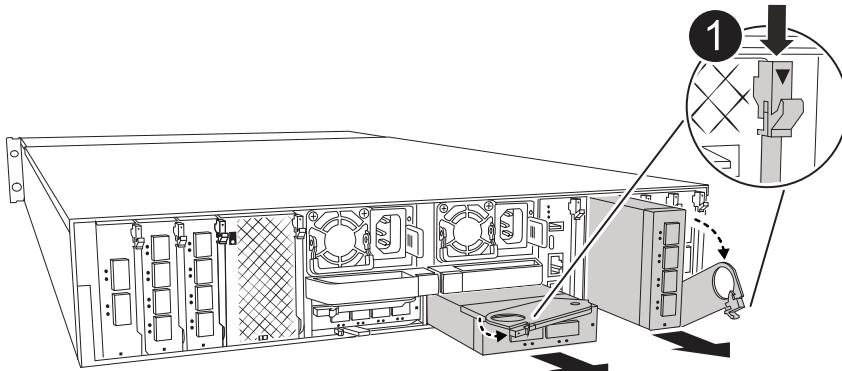


Cables should have been labeled at the beginning of this procedure.

- b. Remove the cable management device from the chassis and set it aside.

Step 2: Remove the I/O cards, NVRAM12, and system management module

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



- | | |
|---|---------------|
| 1 | I/O cam latch |
|---|---------------|

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

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

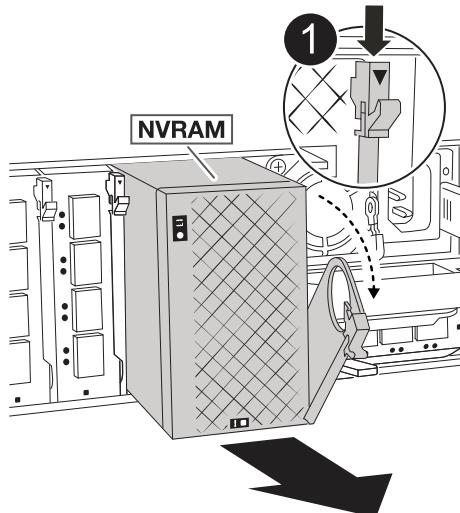
- d. Set the I/O module aside and repeat these steps for any other I/O modules.

2. Remove the NVRAM12 module:

- a. Depress the locking cam button.

The cam button moves away from the chassis.

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

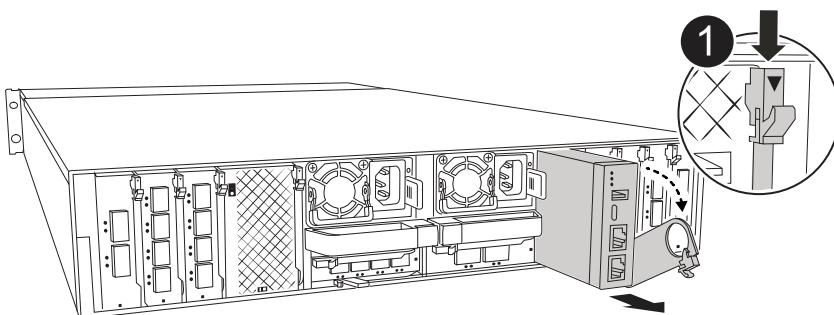


1	NVRAM12 cam latch
---	-------------------

d. Set the NVRAM module on a stable surface.

3. Remove the system management module:

- Depress the cam button on the System Management module.
- Rotate the cam lever down as far as it will go.
- Loop your finger into the hole on the cam lever and pull the module straight out of the system.

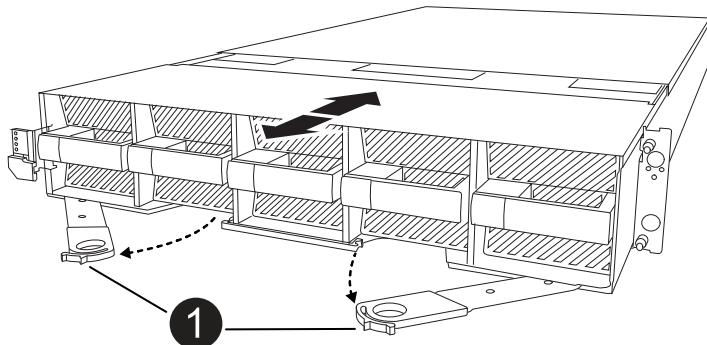


1	System Management module cam latch
---	------------------------------------

Step 3: Remove the controller module

- On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

The controller module moves slightly out of the chassis.



1

Locking cam latches

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

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

Step 4: Replace the impaired chassis

Remove the impaired chassis and install the replacement chassis.

Steps

1. Remove the impaired chassis:
 - a. Remove the screws from the chassis mount points.
 - b. Slide the impaired chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
2. Install the replacement chassis:
 - a. Install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
 - b. Slide the chassis all the way into the equipment rack or system cabinet.
 - c. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the impaired chassis.

Step 5: Install the chassis components

After the replacement chassis is installed, you need to install the controller module, recable the I/O modules and system management module, and then reinstall and plug in the PSUs.

Steps

1. Install the controller module:
 - a. Align the end of the controller module with the opening in the front of the chassis, and then gently push the controller all the way into the chassis.
 - b. Rotate the locking latches into the locked position.
2. Install the I/O cards at the rear of the chassis:
 - a. Align the end of the I/O module with the same slot in the replacement chassis as in the impaired chassis, and then gently push the module all the way into the chassis.

- b. Rotate the cam latch upward into the locked position.
 - c. Repeat these steps for any other I/O modules.
3. Install the system management module at the rear of the chassis:
 - a. Align the end of the system management module with the opening in the chassis, and then gently push the module all the way into the chassis.
 - b. Rotate the cam latch upward into the locked position.
 - c. If you have not already done so, reinstall the cable management device and reconnect the cables to the I/O cards and system management module.



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

- Make sure that the cables are connected according to the cable labels.
4. Install the NVRAM12 module in the back of the chassis at the rear of the chassis:
 - a. Align the end of the NVRAM12 module with the opening in the chassis, and then gently push the module all the way into the chassis.
 - b. Rotate the cam latch upward into the locked position.
 5. Install the PSUs:
 - a. Using both hands, support and align the edges of the PSU with the opening in the chassis.
 - b. Gently push the PSU into the chassis until the locking tab clicks into place.

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



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

6. Reconnect the PSU power cables to both of the PSUs and secure each power cable to the PSU using the power cable retainer.

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

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

What's next?

After you've replaced the impaired FAS70 and FAS90 chassis and reinstalled the components into it, you need to [complete the chassis replacement](#).

Complete the chassis replacement - FAS70 and FAS90

Reboot the controller, verify system health, and return the failed part to NetApp to complete the final step in the FAS70 and FAS90 chassis replacement procedure.

Step 1: Boot the controllers and verify system health

After the controllers reboot, boot ONTAP, give back the controllers, and verify the storage system health.

Steps

1. Check the console output:
 - a. If the controller boots to the LOADER prompt, reboot the controller with the `boot_ontap` command.
 - b. If the console displays `waiting for giveback` after the reboot, log into the partner controller and check that the replaced controller is ready for giveback with the `storage failover show` command.
2. Perform the giveback:
 - a. Connect the console cable to the partner controller.
 - b. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
 - c. If automatic giveback was disabled, reenable it: `storage failover modify -node impaired_node_name -auto-giveback true`
 - d. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`
3. After the giveback is complete, run [Active IQ Config Advisor](#) to verify the health of the storage system and correct any issues you encounter.

Step 2: Return the failed part to NetApp

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

Controller

Controller replacement workflow - FAS70 and FAS90

Get started with replacing the controller in your FAS70 or FAS90 storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.

1

[Review controller replacement requirements](#)

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

2

[Shut down the impaired controller](#)

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

3

[Replace the controller](#)

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

4

Restore and verify the system configuration

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

5

Give back the controller

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

6

Complete controller replacement

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

Requirements to replace the controller - FAS70 and FAS90

Before replacing the controller in your FAS70 or FAS90 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

Review the requirements for replacing the controller.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- Do not use this procedure for controller upgrades; instead, refer to the [Choose your controller hardware upgrade procedure](#) for guidance.
- If your system is in a MetroCluster configuration, you must review the section [Choosing the correct recovery procedure](#) to determine whether you should use this controller replacement procedure.
- You must replace the failed component with the field-replaceable unit (FRU) you received from NetApp.
- You must replace a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- Because the boot device is located on the System Management module that is installed in the back of the system, you do not need to move the boot device when replacing a controller module.
- It is important that you apply the commands in these steps on the correct systems:
 - The *impaired* controller is the controller that is being replaced.
 - The *replacement* controller is the new controller that is replacing the impaired controller.
 - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text log file.

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

What's next?

After you've reviewed the requirements to replace your FAS70 or FAS90 controller, you need to [shut down the impaired controller](#).

Shut down the impaired controller - FAS70 and FAS90

Shut down the controller in your FAS70 or FAS90 storage system to prevent data loss and ensure system stability when replacing the controller.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: Controller is in a MetroCluster

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

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

```
metrocluster node show
```

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

If the impaired controller is displaying...	Then...
System prompt or password prompt (enter system password)	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

What's next?

After you've shut down the controller, you need to [replace the controller](#).

Replace the controller - FAS70 and FAS90

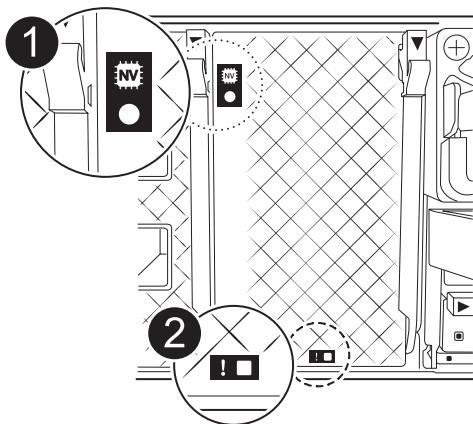
Replace the controller in your FAS70 or FAS90 system when a hardware failure requires it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

Step 1: Remove the controller module

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

Steps

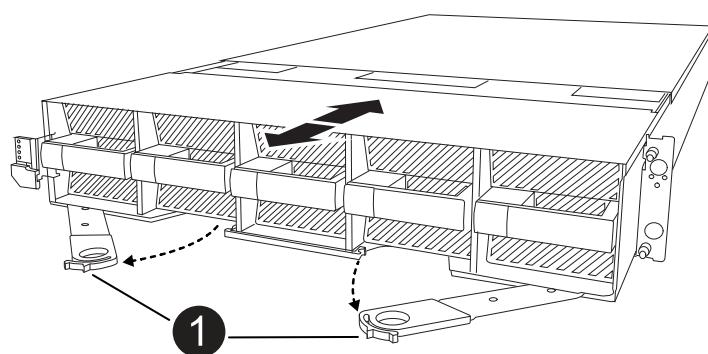
1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:



1	NVRAM status LED
2	NVRAM attention LED

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

The controller module moves slightly out of the enclosure.



1

Locking cam latches

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

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

Step 2: Move the fans

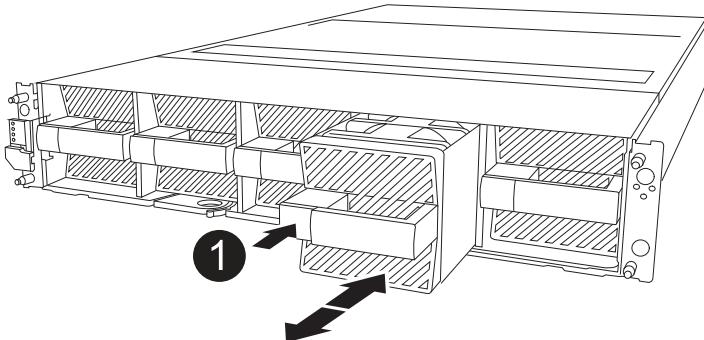
You must remove the five fan modules from the impaired controller module to the replacement controller module.

Steps

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



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



1

Black locking button

4. Install the fan in the replacement controller module:

- Align the edges of the fan housing with the opening in the front of the replacement controller module.
- Gently slide the fan module all the way into the replacement controller module until it locks in place.

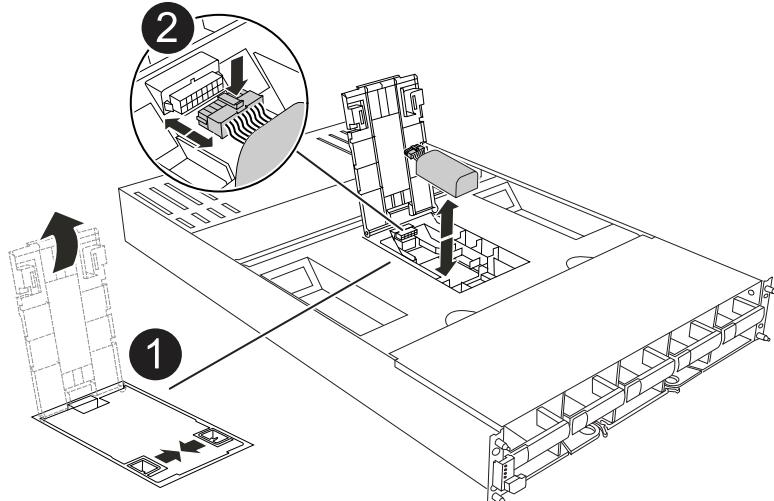
5. Repeat the preceding steps for the remaining fan modules.

Step 3: Move the NV battery

Move the NV battery to the replacement controller.

Steps

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



1

NV battery air duct cover

2

NV battery plug

3

NV battery pack

- Lift the battery up to access the battery plug.

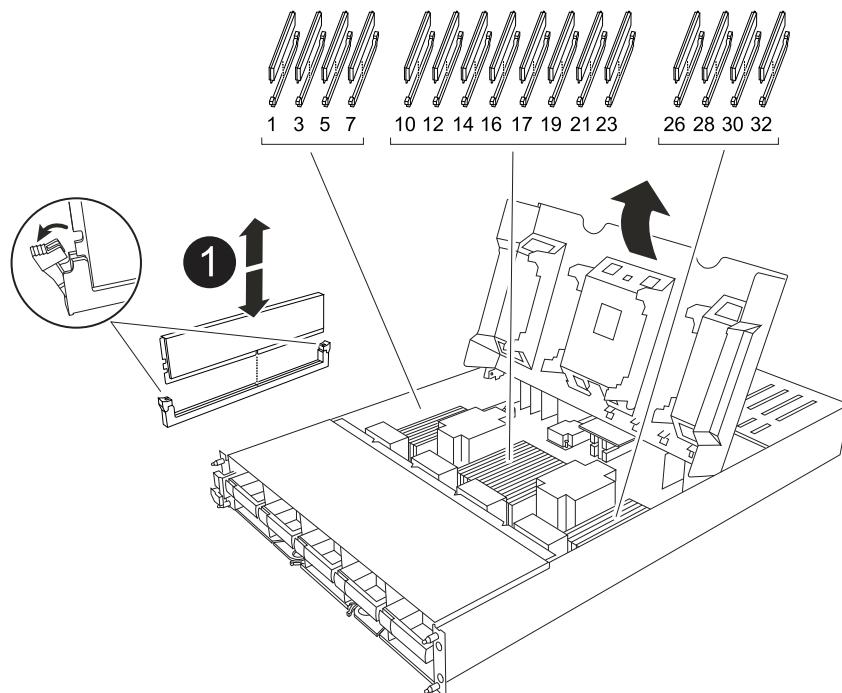
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module.
5. Move the battery pack to the replacement controller module and then install it in the NV battery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

Step 4: Move system DIMMs

Move the DIMMs to the replacement controller module.

Steps

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



1

System DIMM

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



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

4. Locate the slot where you are installing the DIMM in the replacement controller module.

5. Insert the DIMM squarely into the slot.

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



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

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

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

Step 5: Install the controller module

Reinstall the controller module and boot it.

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

It must lie flush against the controller module sheet metal.

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

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



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



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

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



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

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

6. Recable the storage system, as needed.

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

What's next?

After you've replaced the impaired FAS70 or FAS90 controller, you need to [restore the system configuration](#).

Restore and verify the system configuration - FAS70 and FAS90

Verify that the controller's HA configuration is active and functioning correctly in your FAS70 or FAS90 storage system, and confirm that the system's adapters list all the paths to the disks.

Step 1: Verify HA config settings

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

Steps

1. Boot to maintenance mode: `boot_ontap maint`

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

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

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



If you see *PERSONALITY MISMATCH* contact customer support.

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

4. Verify that all components display the same HA state: `ha-config show`

The HA state should be the same for all components.

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

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

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

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

Step 2: Verify disk list

Steps

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

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

2. Exit Maintenance mode: `halt`.

What's next?

After you've restored and verified the system configuration for your FAS70 or FAS90 system, you need to [give back the controller](#).

Give back the controller - FAS70 and FAS90

Return control of storage resources to the replacement controller so your FAS70 or FAS90 system can resume normal operation. The give back procedure varies based on

the encryption type used by your system: no encryption or Onboard Key Manager (OKM) encryption.

No encryption

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

Steps

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

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

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



You are prompted twice for the passphrase.

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



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

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

The output should show no results when when properly synchronized.

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

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

What's next?

After you've transferred the ownership of storage resources back to the replacement controller, you need to [complete the controller replacement](#) procedure.

Complete controller replacement - FAS70 and FAS90

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

Step 1: Verify LIFs and check cluster health

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

Steps

1. Verify that the logical interfaces are reporting to their home server and ports:

```
network interface show -is-home false
```

If any logical interfaces are listed as false, revert them to their home ports:

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

2. Check the health of your cluster. See the [How to perform a cluster health check with a script in ONTAP](#) KB article.
3. If automatic giveback was disabled, reenable it:

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

Step 2: Return the failed part to NetApp

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

Replace a DIMM - FAS70 and FAS90

Replace a DIMM in your FAS70 or FAS90 system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

Before you begin

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

3. Take the impaired controller to the LOADER prompt:

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

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: Controller is in a MetroCluster

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

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

```
metrocluster node show
```

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

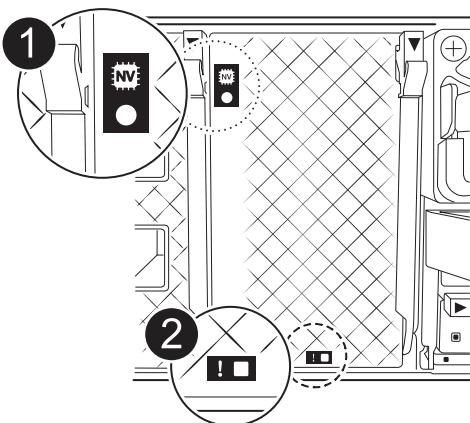
If the impaired controller is displaying...	Then...
System prompt or password prompt (enter system password)	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Step 2: Remove the controller module

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

Steps

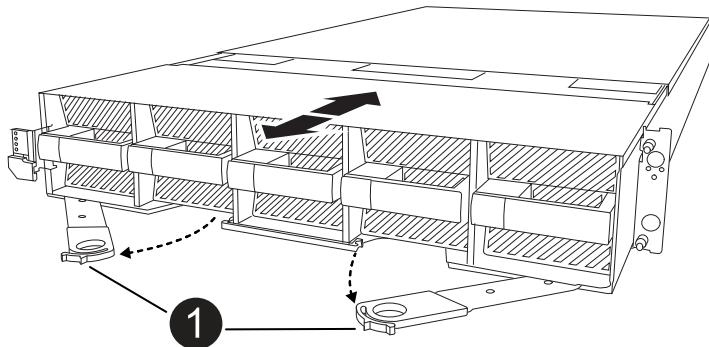
1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:



1	NVRAM status LED
2	NVRAM attention LED

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

The controller module moves slightly out of the enclosure.



1

Locking cam latches

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

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

Step 3: Replace a DIMM

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

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

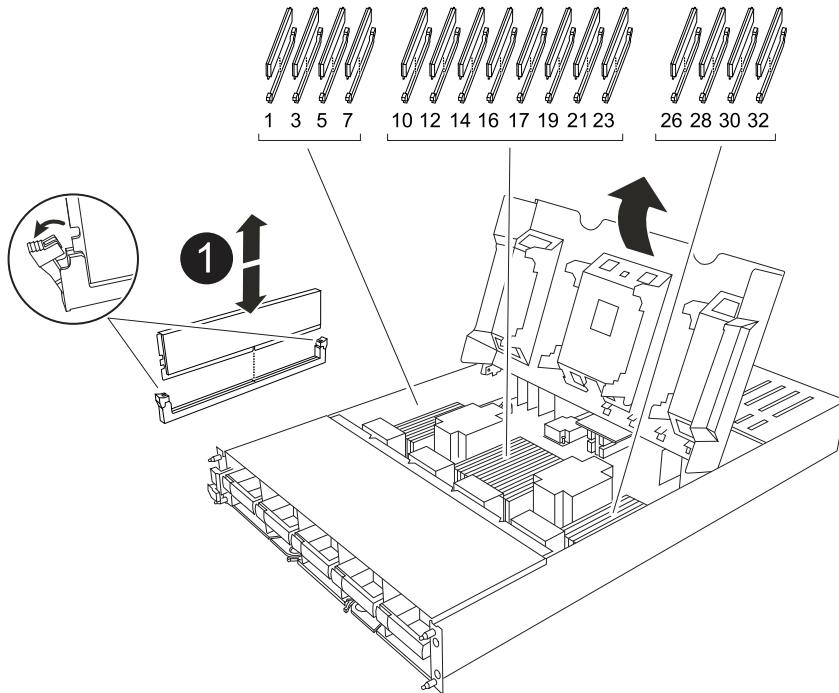
DIMMs locations are dependent on the system model:

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

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



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



1

DIMM and DIMM ejector tabs

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

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

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

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



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

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

Step 4: Install the controller

Reinstall the controller module and boot it.

Steps

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

It must lie flush against the controller module sheet metal.

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

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



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

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

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

Step 5: Return the failed part to NetApp

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

Replace a fan - FAS70 and FAS90

Replace a fan module in your FAS70 or FAS90 system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

Steps

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

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

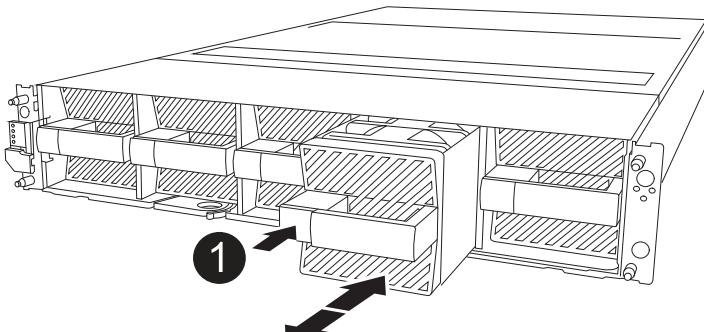


There is a single LED for each fan. It is green when the fan is functioning correctly and amber when not.

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



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



1	Black release button
----------	----------------------

5. Set the fan module aside.
6. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.
When inserted into a live system, the amber Attention LED turns off once the fan is recognized by that system.
7. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

Replace the Flash Cache module carrier or a caching module - FAS70 and FAS90

The NVMe SSD Flash Cache module carrier in your FAS70 or FAS90 system contains one or two Flash Cache modules (caching modules) with a single SSD Flash Cache drive integrated into each caching module.

The FAS70 supports 2TB caching modules and FAS90 supports 4TB caching modules. You cannot mix caching modules of different capacity in the Flash Cache module carrier.

You can perform either of the following procedures depending on what component you need to replace: the entire Flash Cache module carrier or a caching module.

- [Replace the Flash Cache module carrier](#)
- [Replace the caching module](#)

Replace the Flash Cache module carrier

The Flash Cache module carrier is located in slot 6 and houses up to two Flash Cache modules. You cannot hot-swap the Flash Cache module carrier.

Before you begin

- Ensure your storage system has the appropriate operating system for the replacement Flash Cache module carrier.
- Confirm all other components are functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired node

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

3. Take the impaired controller to the LOADER prompt:

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

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: Controller is in a MetroCluster

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

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

```
metrocluster node show
```

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

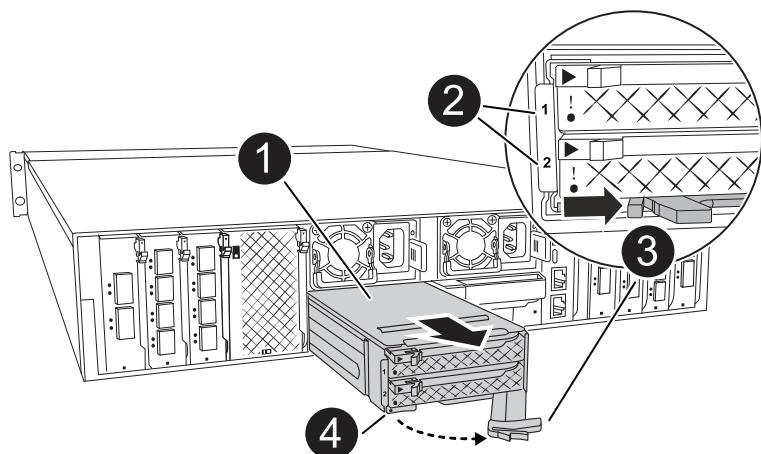
If the impaired controller is displaying...	Then...
System prompt or password prompt (enter system password)	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Step 2: Replace the Flash Cache module carrier

Perform the following steps to replace the Flash Cache module carrier.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the failed Flash Cache module carrier, in slot 6, by the lit amber Attention LED on the front of the Flash Cache module carrier.



1	Flash Cache module carrier
2	Caching module slot numbers
3	Flash Cache module carrier cam handle
4	Flash Cache module carrier fault LED

3. Remove the failed Flash Cache module carrier:

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

- b. Pinch the blue tab at the bottom of the Flash Cache module carrier.
 - c. Rotate the tab away from the module.
4. Pull the Flash Cache module carrier out of the controller module and set it on an antistatic mat.
5. Move the caching modules to the replacement Flash Cache module carrier:
 - a. Pinch the Terra Cotta tab at the top of the caching module and rotate the cam handle away from the caching module.
 - b. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the Flash Cache module carrier.
 - c. Install the caching module into the same slot in the replacement Flash Cache module carrier and rotate the cam handle to the closed position on the caching module to lock it in place.
6. Repeat these steps if there is a second caching module.
7. Install the replacement Flash Cache module carrier into the system:
 - a. Align the module with the edges of the enclosure slot opening.
 - b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
 - c. Rotate the cable management tray up to the closed position.

Step 3: Reboot the controller

After you replace the Flash Cache module carrier, you must reboot the controller module.

Steps

1. From the LOADER prompt, reboot the node: *bye*



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

2. Return the node to normal operation: *storage failover giveback -ofnode impaired_node_name*
3. If automatic giveback was disabled, reenable it: *storage failover modify -node local -auto-giveback true*

Step 4: Return the failed part to NetApp

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

Replace the caching module

The Flash Cache modules (caching modules) are located in slot 6-1 or in slot 6-2 or in both slot 6-1 and slot 6-2.

You can hot-swap the individual caching modules with caching modules of the same capacity from the same vendor or from a different supported vendor.

Before you begin

- Ensure the replacement caching module has the same capacity as the failed one, from the same vendor or from a different supported vendor.
- Confirm all other components are functioning properly; if not, you must contact technical support.
- The drives in the caching modules are not field replaceable units (FRU). You must replace the entire

caching module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the failed caching module, in slot 6, by the lit amber Attention LED on the front of the caching module.
3. Prepare the caching module slot for replacement as follows:
 - a. Record the caching module capacity, part number, and serial number on the target node: `system node run local sysconfig -av 6`
 - b. In admin privilege level, prepare the target caching module slot for removal, responding `y` when prompted whether to continue: `system controller slot module remove -node node_name -slot slot_number` The following command prepares slot 6-1 on node1 for removal, and displays a message that it is safe to remove:

```
::> system controller slot module remove -node node1 -slot 6-1
```

Warning: SSD module in slot 6-1 of the node node1 will be powered off for removal.

Do you want to continue? (y|n): y

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

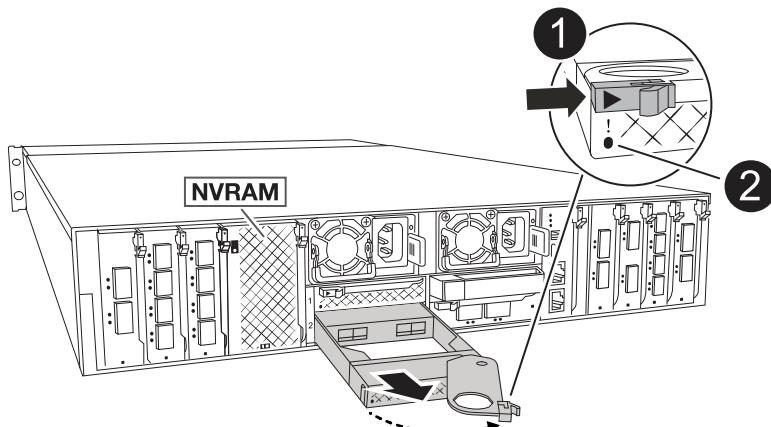
- c. Display the slot status with the `system controller slot module show` command.

The caching module slot status displays `powered-off` in the screen output for the caching module that needs replacing.



See the [Command man pages](#) for your version of ONTAP for more details.

4. Remove the caching module:



1

Caching module cam handle

2

Caching module fault LED

- a. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
- b. Press the terra cotta release button on the front of the caching module.
- c. Rotate the cam handle as far as it will go.
- d. Remove the caching module module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the Flash Cache module carrier.

Be sure to support the caching module as you remove it from the Flash Cache module carrier.

5. Install the replacement caching module:

- a. Align the edges of the caching module with the opening in the controller module.
- b. Gently push the caching module into the bay until the cam handle engages.
- c. Rotate the cam handle until it locks into place.
- d. Rotate the cable management tray up to the closed position.

6. Bring the replacement caching module online by using the `system controller slot module insert` command as follows:

The following command prepares slot 6-1 on node1 for power-on, and displays a message that it is powered on:

```
::> system controller slot module insert -node node1 -slot 6-1  
Warning: NVMe module in slot 6-1 of the node localhost will be powered  
on and initialized.  
Do you want to continue? (y|n): `y`
```

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

7. Verify the slot status using the `system controller slot module show` command.

Make sure that command output reports status for the as powered-on and ready for operation.

8. Verify that the replacement caching module is online and recognized, and then visually confirm that the amber attention LED is not lit: `sysconfig -av slot_number`



If you replace the caching module with a caching module from a different vendor, the new vendor name is displayed in the command output.

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

Replace NVRAM - FAS70 and FAS90

Replace the NVRAM in your FAS70 or FAS90 system when the non-volatile memory becomes faulty or requires an upgrade. The replacement process involves shutting down the impaired controller, replacing the NVRAM module or the NVRAM DIMM, reassigning the disks, and returning the failed part to NetApp.

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

Before you begin

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

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

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: Controller is in a MetroCluster

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

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

```
metrocluster node show
```

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

If the impaired controller is displaying...	Then...
System prompt or password prompt (enter system password)	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Step 2: Replace the NVRAM module or NVRAM DIMM

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

Option 1: Replace the NVRAM module

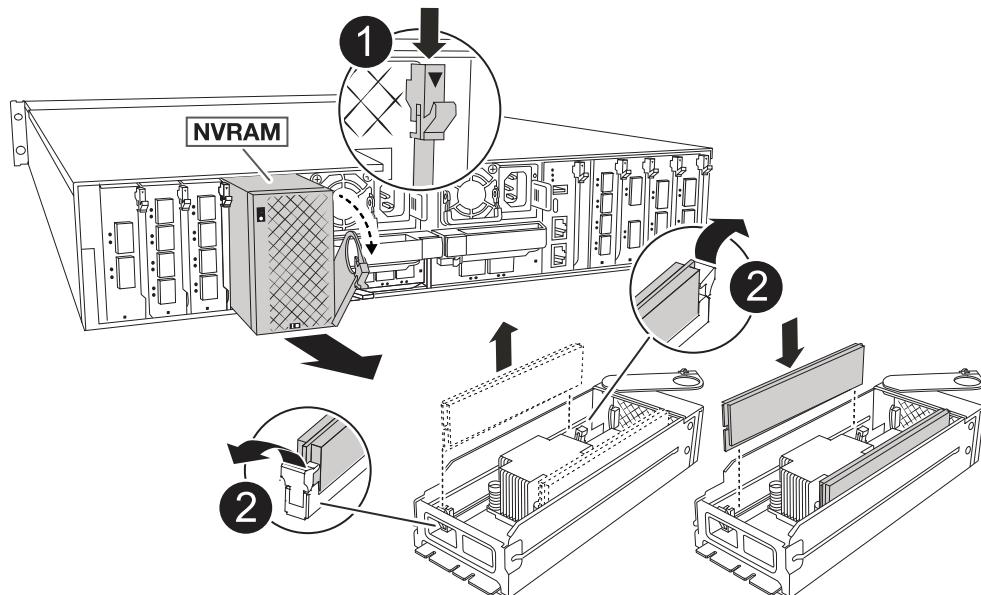
To replace the NVRAM module, locate it in slot 4/5 in the enclosure and follow the specific sequence of steps.

Steps

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

The cam button moves away from the enclosure.

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



1	Cam locking button
2	DIMM locking tabs

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

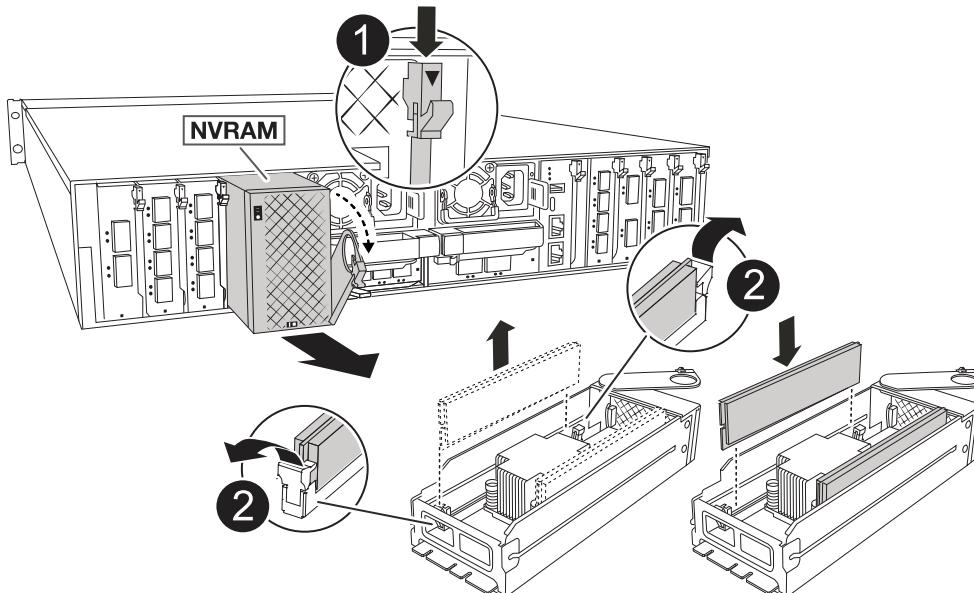
- b. Gently slide the module into the slot all the way, and then rotate the cam latch all the way up to lock the module in place.
8. Recable the controller.
9. Rotate the cable management tray up to the closed position.

Option 2: Replace the NVRAM DIMM

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

Steps

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



1	Cam locking button
2	DIMM locking tabs

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



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

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

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

Step 3: Reboot the controller

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

Steps

1. Plug the power cables back into the PSU.

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

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

Step 4: Reassign disks

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



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

Steps

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

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

```

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

```

4. Give back the controller:

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

The controller takes back its storage and completes booting.

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



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

For more information, see the [Manual giveback commands](#) topic to override the veto.

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

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

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

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

```

node1:> storage disk show -ownership

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

```

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

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

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

This is required if both of the following are true:

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

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

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

```

node1_siteA:> metrocluster node show -fields configuration-state

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

4 entries were displayed.

```

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

Step 5: Return the failed part to NetApp

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

Replace the NV battery - FAS70 and FAS90

Replace the NV battery in your FAS70 or FAS90 system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: Controller is in a MetroCluster

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

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

```
metrocluster node show
```

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

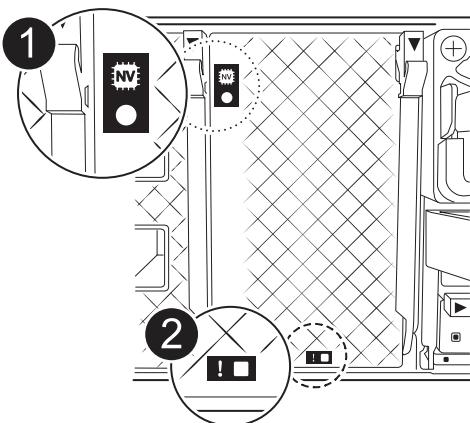
If the impaired controller is displaying...	Then...
System prompt or password prompt (enter system password)	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Step 2: Remove the controller module

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

Steps

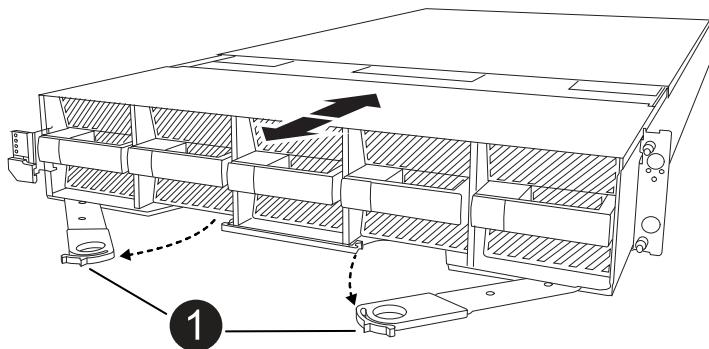
1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:



1	NVRAM status LED
2	NVRAM attention LED

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

The controller module moves slightly out of the enclosure.



1	Locking cam latches
---	---------------------

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

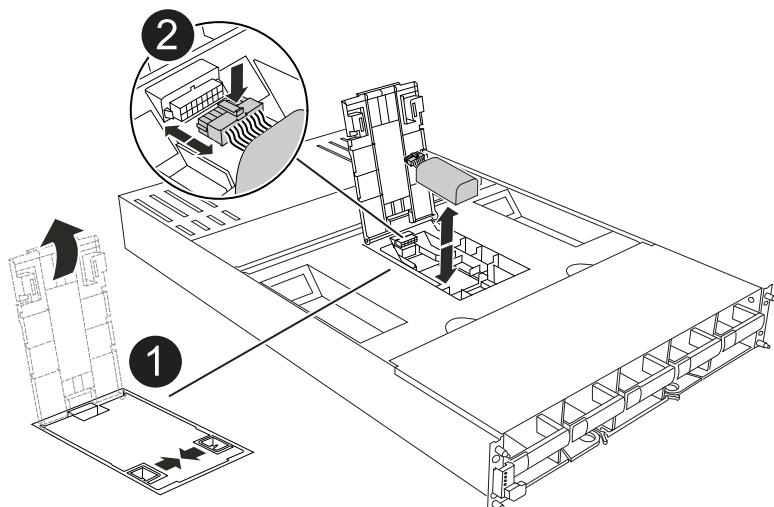
Make sure that you support the bottom of the controller module as you slide it out of the enclosure.

Step 3: Replace the NV battery

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

Steps

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



1	NV battery air duct cover
2	NV battery plug

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

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

Make sure that the plug locks into the socket.

Step 4: Reinstall the controller module

Reinstall the controller module and boot it.

Steps

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

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



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

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

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

Step 5: Return the failed part to NetApp

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

I/O module

Overview of add and replace an I/O module - FAS70 and FAS90

The FAS70 or FAS90 system offers flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your FAS70 or FAS90 storage system with the same type of I/O module,

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

- [Add an I/O module](#)

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

- [Hot-swap an I/O module](#)

You can hot-swap certain I/O modules for an equivalent I/O module to restore the storage system to its optimal operating state. Hot-swapping is done without having to perform a manual takeover.

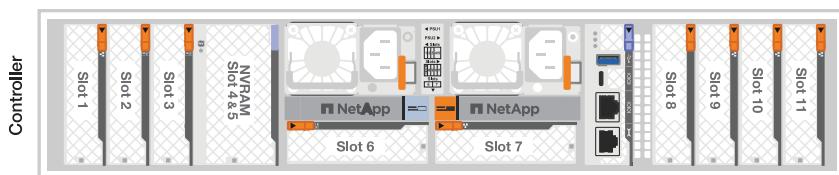
To use this procedure, your storage system must be running ONTAP 9.18.1 or later.

- [Replace an I/O module](#)

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

I/O slot numbering

The I/O slots on FAS70 and FAS90 controllers are numbered 1 through 11, as shown in the following illustration.



Add an I/O module - FAS70 and FAS90

Add an I/O module to your FAS70 and FAS90 system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your FAS70 and FAS90 storage system when there are empty slots available or when all slots are fully populated.

Step 1: Shut down the impaired controller module

Shut down or take over the impaired controller module.

Option 1: Most systems

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

Before you begin

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

Steps

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

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

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

2. Disable automatic giveback from the console of the impaired controller:

```
storage failover modify -node impaired-node -auto-giveback-of false
```



When you see *Do you want to disable auto-giveback?*, enter `y`.

- a. If the impaired controller cannot be brought up or is already taken over, you must take the HA interconnect link down from the healthy controller before booting up the impaired controller. This will prevent the impaired controller from performing automatic giveback.

```
system ha interconnect link off -node healthy-node -link 0
```

```
system ha interconnect link off -node healthy-node -link 1
```

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
System prompt or password prompt (enter system password)	<p>Halt or take over the impaired controller from the healthy controller: <code>storage failover takeover -ofnode impaired_node_name</code></p> <p>When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <code>y</code>.</p>

Option 2: MetroCluster configuration



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

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

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

Steps

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

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

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

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

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

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next Step.
Waiting for giveback...	Press Ctrl-C, and then respond <code>y</code> when prompted.
System prompt or password prompt (enter system password)	<p>Halt or take over the impaired controller from the healthy controller: <code>storage failover takeover -ofnode impaired_node_name</code></p> <p>When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <code>y</code>.</p>

Step 2: Add the new I/O module

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

Before you begin

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

- Make sure that all other components are functioning properly.
- Make sure you have the replacement component you received from NetApp.

Add I/O module to an available slot

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

Steps

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



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

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

7. From the LOADER prompt, reboot the node:

bye



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

8. Give back the controller from the partner controller:

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

9. Repeat these steps for controller B.

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

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

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

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

Add I/O module to a fully-populated system

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

About this task

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

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

Steps

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

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

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

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



If you encounter an issue during reboot, see [BURT 1494308 - Environment shutdown might be triggered during I/O module replacement](#)

10. Give back the controller from the partner controller:

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

11. Enable automatic giveback if it was disabled:

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

12. Do one of the following:

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

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

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

13. Repeat these steps for controller B.

Hot swap an I/O module - FAS 70 and FAS 90

You can hot swap an Ethernet I/O module in your FAS 70 or FAS 90 storage system if a module fails and your storage system meets all ONTAP version requirements.

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

About this task

- You do not need to perform a manual takeover before replacing the failed I/O module.
- Apply commands to the correct controller and I/O slot during the hot-swap:
 - The *impaired controller* is the controller where you are replacing the I/O module.
 - The *healthy controller* is the HA partner of the impaired controller.
- You can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

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

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

Step 1: Ensure the storage system meets the procedure requirements

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



If your storage system is not running ONTAP 9.18.1 GA or later, you cannot use this procedure, you must use the [replace an I/O module procedure](#).

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

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

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

supported for your storage system.

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

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

- The controllers in your storage system can be in either of the following states:
 - Both controllers can be up and running I/O (serving data).
 - Either controller can be in a takeover state if the takeover was caused by the failed I/O module and the nodes are otherwise functioning properly.

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

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

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

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

Steps

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



The I/O module should be failed (ports should be in the link down state); however, if the links are still up and they contain the last functioning cluster port, unplugging the cables triggers an automatic takeover.

Wait five minutes after unplugging the cables to ensure any takeovers or LIF failovers complete before continuing with this procedure.

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

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

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

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

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

If...	Then...
If either controller took over its partner automatically	<p>Disable automatic giveback:</p> <ol style="list-style-type: none"> 1. Enter the following command from the console of the controller that took over its partner: <pre>storage failover modify -node local -auto -giveback false</pre> <ol style="list-style-type: none"> 2. Enter <i>y</i> when you see the prompt <i>Do you want to disable auto-giveback?</i>
Both controllers are up and running I/O (serving data)	Go to the next step.

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

- a. Enter the following command:

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

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

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

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

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

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

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

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

```
system controller slot module show
```

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

Step 3: Replace the failed I/O module

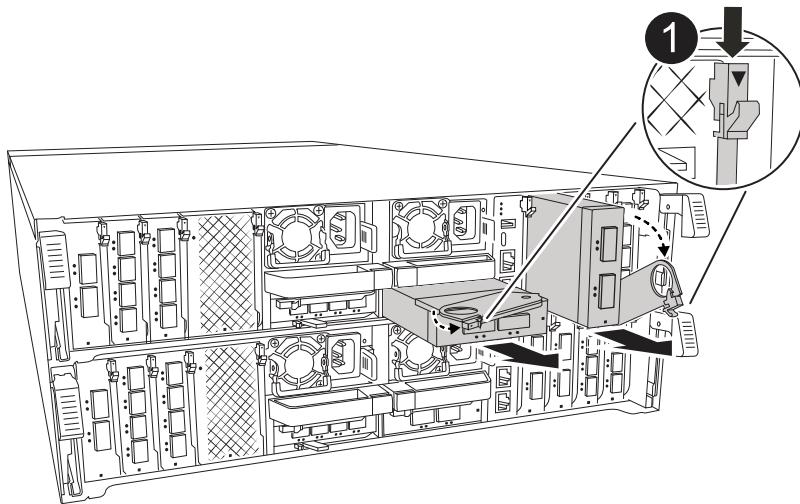
Replace the failed I/O module with an equivalent I/O module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
3. Remove the I/O module from the controller module:



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



1	Cam locking button
----------	--------------------

- a. Depress the cam latch button.
- b. Rotate the cam latch away from the module as far as it will go.
- c. Remove the module from the controller module by hooking your finger into the cam lever opening and pulling the module out of the controller module.

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

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

Step 4: Bring the replacement I/O module online

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

About this task

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

I/O module.

Steps

1. Bring the replacement I/O module online:

- a. Enter the following command:

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

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

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

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

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

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

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

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

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

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

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



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

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

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

```

node2::> event log show -event *hotplug.init*
Time           Node           Severity      Event
-----
-----
7/11/2025 16:04:06  node2      NOTICE       hotplug.init.success:
Initialization of ports "e7b" in slot 7 succeeded

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

2 entries were displayed.

```

- b. If the port initialization fails, review the EMS log for the next steps to take.
3. Verify the I/O module slot is powered on and ready for operation:

```
system controller slot module show
```

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

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

Enter the command from the console of the impaired controller:

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

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

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

```

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

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

```

Step 5: Restore the storage system to normal operation

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

Steps

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

If...	Then...
If either controller took over its partner automatically	<ol style="list-style-type: none"> 1. Return the controller that was taken over to normal operation by giving back its storage: <code>storage failover giveback -ofnode controller that was taken over_name</code> 2. Restore automatic giveback from the console of the controller that was taken over: <code>storage failover modify -node local -auto -giveback true</code>

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

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

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

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

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

Step 6: Return the failed part to NetApp

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

Replace an I/O module - FAS70 and FAS90

Replace an I/O module in your FAS70 or FAS90 system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

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

Before you begin

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

Step 1: Shut down the impaired node

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

3. Take the impaired controller to the LOADER prompt:

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

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: Controller is in a MetroCluster



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

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

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

Steps

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

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

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

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

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

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

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

Step 2: Replace a failed I/O module

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

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.

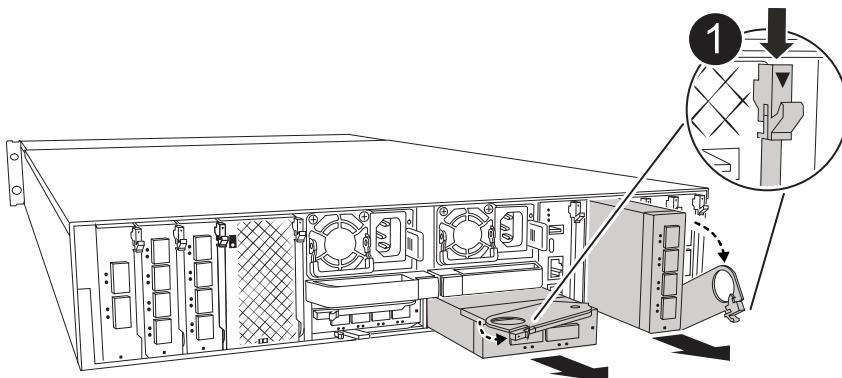


Make sure that you label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.

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



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



1

I/O cam latch

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

4. Remove the target I/O module from the enclosure:
 - a. Depress the cam button on the target module.
 - b. Rotate the cam latch away from the module as far as it will go.

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

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

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

Step 3: Reboot the controller

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

Steps

1. Reboot the controller from the LOADER prompt:

bye



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

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

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

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

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

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

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

Step 4: Return the failed part to NetApp

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

Hot-swap a power supply - FAS70 and FAS90

Replace an AC or DC power supply unit (PSU) in your FAS70 or FAS90 system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cable, replacing the faulty PSU, and then reconnecting it to the power source.

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

a PSU.

About this task

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



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

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

Option 1: Hot-swap an AC PSU

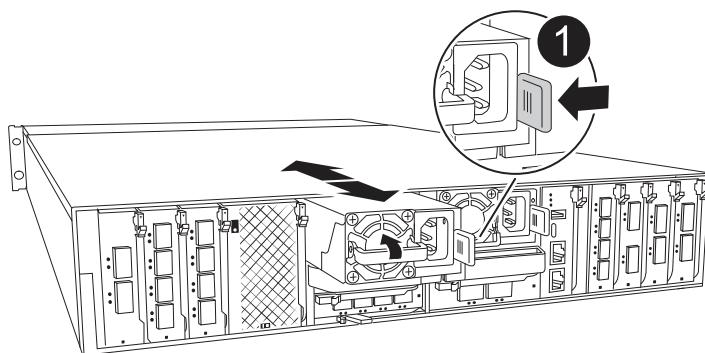
To replace an AC PSU, complete the following steps.

Steps

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



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



Terracotta PSU locking tab

5. Install the replacement PSU in the controller module:

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

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



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

6. Reconnect the PSU cabling:

- a. Reconnect the power cable to the PSU.
- b. Secure the power cable to the PSU using the power cable retainer.

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

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

Option 2: Hot-swap a DC PSU

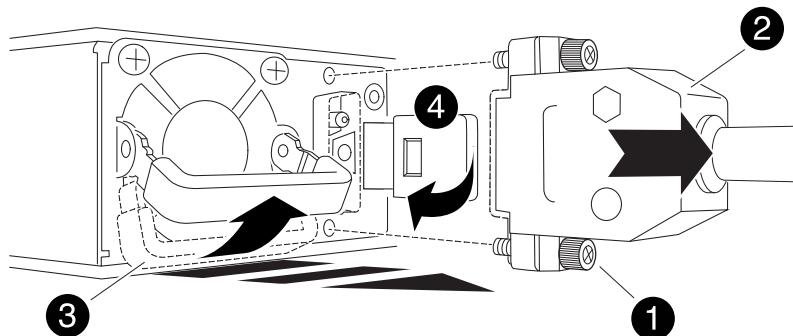
To replace a DC PSU, complete the following steps.

Steps

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



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



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

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

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

way.



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

6. Reconnect the D-SUB DC power cable:

- a. Plug the power cable connector into the PSU.
- b. Secure the power cable to the PSU with the thumbscrews.

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

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

Replace the real-time clock battery - FAS70 and FAS90

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your FAS70 or FAS90 system to ensure that services and applications relying on accurate time synchronization remain operational.

Before you begin

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

You must use an approved RTC battery.

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

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

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: Controller is in a MetroCluster

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

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

```
metrocluster node show
```

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

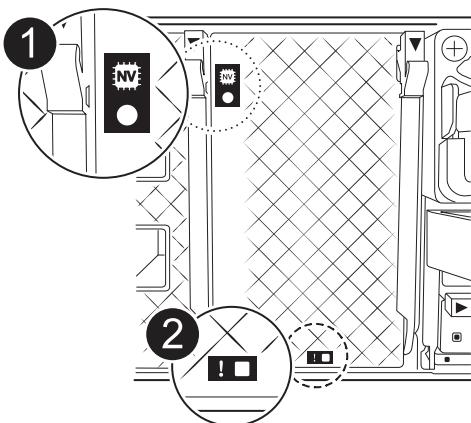
If the impaired controller is displaying...	Then...
System prompt or password prompt (enter system password)	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Step 2: Remove the controller module

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

Steps

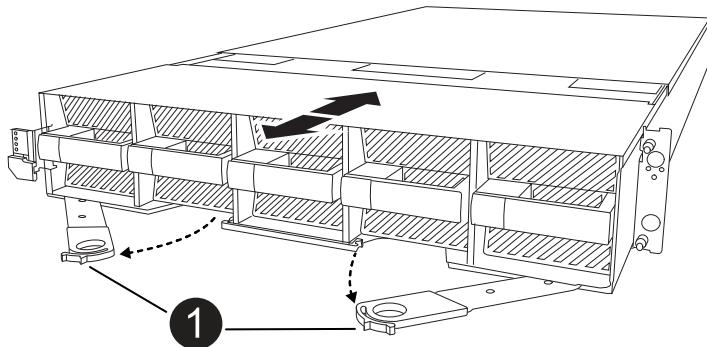
1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:



1	NVRAM status LED
2	NVRAM attention LED

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

The controller module moves slightly out of the enclosure.



1

Locking cam latches

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

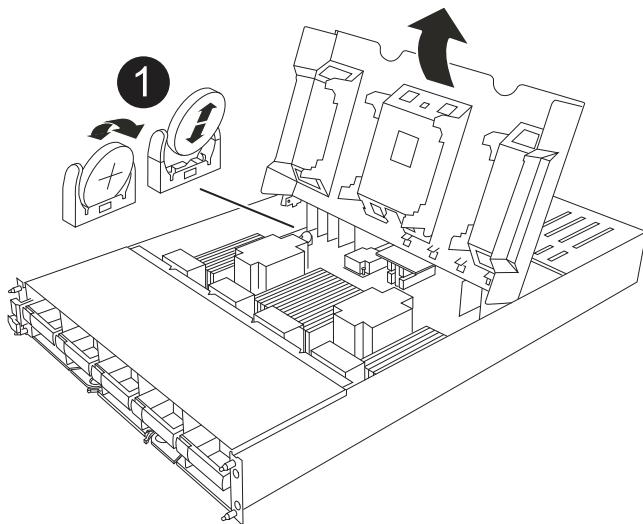
Make sure that you support the bottom of the controller module as you slide it out of the enclosure.

Step 3: Replace the RTC battery

Remove failed RTC battery and install the replacement RTC battery.

Steps

1. Open the controller air duct on the top of the controller.
 - a. Insert your fingers in the recesses at the far ends of the air duct.
 - b. Lift the air duct and rotate it upward as far as it will go.
2. Locate the RTC battery under the air duct.



1

RTC battery and housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.



Observe the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder indicates the correct orientation.

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

Step 4: Reinstall the controller module

Reinstall the controller module and boot it.

Steps

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

It must lie flush against the controller module sheet metal.

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



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

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

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

Step 5: Reset the time and date on the controller



After replacing the RTC battery, inserting controller and powering on first BIOS reset, you will see the following error messages: RTC date/time error. Reset date/time to default RTC power failure error. These messages are expected and you can continue with this procedure.

Steps

1. Check the date and time on the healthy controller with the `cluster date show` command.



If your system stops at the boot menu, select the option for `Reboot node` and respond `y` when prompted, then boot to LOADER by pressing `Ctrl-C`

- a. At the LOADER prompt on the target controller, check the time and date with the `cluster date show` command.
- b. If necessary, modify the date with the `set date mm/dd/yyyy` command.
- c. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
 1. Confirm the date and time on the target controller.
 2. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.

Step 6: Return the failed part to NetApp

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

Replace system management module - FAS70 and FAS90

Replace the System Management module in your FAS70 or FAS90 system when it becomes defective or its firmware is corrupted. The replacement process involves shutting down the controller, replacing the failed System Management module, rebooting the controller, updating the license keys, and returning the failed part to NetApp.

The System Management module, located at the back of the controller in slot 8, contains onboard components for system management, as well as ports for external management. The target controller must be shut down to replace an impaired System Management module or replace the boot media.

The System Management module has the following onboard components:

- Boot media, allowing boot media replacement without removing the controller module.
- BMC
- Management switch

The System Management module also contains the following ports for external management:

- RJ45 Serial
- USB Serial (Type-C)
- USB Type-A (Boot recovery)
- e0M RJ45 Ethernet

Before you begin

- Make sure all other system components are working properly.
- Make sure that the partner controller is able to take over the impaired controller.
- Make sure you replace the failed component with a replacement component you received from NetApp.

About this task

This procedure uses the following terminology:

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

```
metrocluster node show
```

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message
MAINT=number_of_hours_downh
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:*> system node autosupport invoke -node * -type all -message
MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback false
```

- b. Enter *y* when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next section.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.

If the impaired controller is displaying...	Then...
System prompt or password prompt (enter system password)	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Step 2: Replace the impaired System Management module

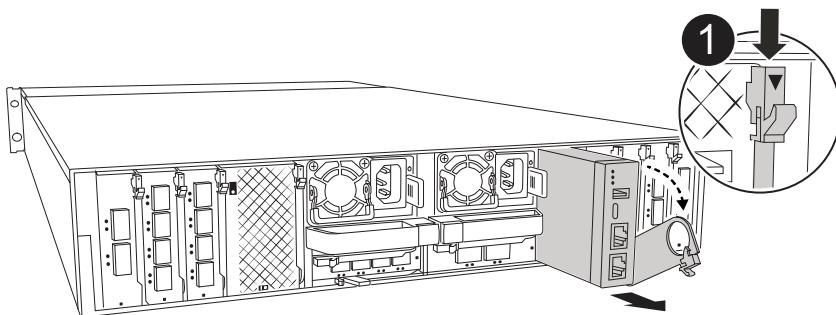
Replace the impaired system management module.

Steps

1. Remove the System Management module:



Make sure NVRAM destage has completed before proceeding. When the LED on the NV module is off, NVRAM is destaged. If the LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

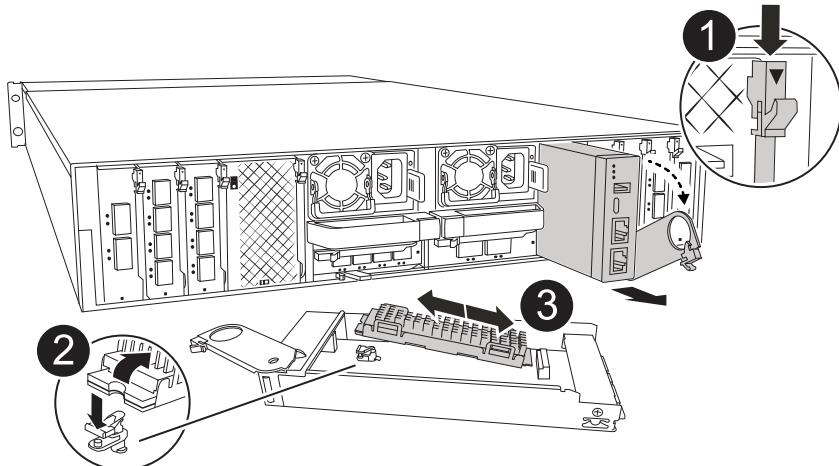


1

System Management module cam latch

- a. If you are not already grounded, properly ground yourself.
- b. Unplug the power supply cables from the PSUs.
2. Remove the System Management module
 - a. Remove any cables connected to the System Management module. Make sure that label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
 - b. Disconnect the power cords from the PSU for the impaired controller.
 - c. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
 - d. Depress the cam button on the System Management module.
 - e. Rotate the cam lever down as far as it will go.

- f. Loop your finger into the hole on the cam lever and pull the module straight out of the system.
 - g. Place the System Management module on an anti-static mat, so that the boot media is accessible.
3. Move the boot media to the replacement System Management module:



1	System Management module cam latch
2	Boot media locking button
3	Boot media

- a. Press the blue boot media locking button in the impaired System Management module.
 - b. Rotate the boot media up and slide it out of the socket.
4. Install the boot media in the replacement System Management module:
- a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
 - b. Rotate the boot media down until it touches the locking button.
 - c. Depress the blue locking and rotate the boot media all the way down and release the blue locking button.
5. Install the replacement System Management module into the enclosure:
- a. Align the edges of the replacement System Management module with the system opening and gently push it into the controller module.
 - b. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.
6. Rotate the cable management arm up to the closed position.
7. Recable the System Management module.

Step 3: Reboot the controller module

Reboot the controller module.

Steps

1. Plug the power cables back into the PSU.

The system will begin to reboot, typically to the LOADER prompt.

2. Enter `bye` at the LOADER prompt.

3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`.

4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.

5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 4: Install licenses and register serial number

You must install new licenses for the node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the node.

However, if the node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed. Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the node as soon as possible.

Before you begin

The license keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

 If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in [Post Motherboard Replacement Process to update Licensing on a AFF/FAS system](#). If you are unsure of the initial ONTAP release for your system, see [NetApp Hardware Universe](#) for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the [NetApp Support Site](#) in the My Support section under Software licenses.



The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`

3. Remove the old licenses, if desired:

- a. Check for unused licenses: `license clean-up -unused -simulate`

- b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

4. Register the system serial number with NetApp Support.
 - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
 - If AutoSupport is not enabled, call [NetApp Support](#) to register the serial number.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

Key specifications

Key specifications for FAS70

The following are select specifications for the FAS70 storage system in a single high availability pair. Visit NetApp Hardware Universe (HWU) for the complete specifications for this storage system.

FAS70 specifications at a glance

- Platform Configuration: FAS70 Dual Chassis HA Pair
- Max Raw Capacity: 15.0000 PB
- Memory: 256.0000 GB
- Form Factor: 2U chassis with 1 HA controllers
- ONTAP Version: ONTAP: 9.16.1P2
- PCIe Expansion Slots: 18
- Minimum ONTAP Version: ONTAP 9.15.1

Scaleout maximums

- Type: NAS; HA Pairs: 12; Raw Capacity: 180.0 PB / 159.9 PiB; Max Memory: 3072 GB
- Type: SAN; HA Pairs: 6; Raw Capacity: 90.0 PB / 79.9 PiB; Max Memory: 1536 GB
- Type: HA Pair; Raw Capacity: 15.0 PB / 13.3 PiB; Max Memory: 256.0000

I/O

Onboard I/O

No onboard I/O data.

Total I/O

- Protocol: Ethernet 200 Gbps; Ports: 20
- Protocol: Ethernet 100 Gbps; Ports: 32
- Protocol: Ethernet 25 Gbps; Ports: 48
- Protocol: Ethernet 10 Gbps; Ports: 48
- Protocol: FC 64 Gbps; Ports: 48

- Protocol: NVMe/FC 64 Gbps; Ports: 48
- Ports: 0
- Protocol: SAS 12 Gbps; Ports: 48

Management ports

- Protocol: Ethernet 1 Gbps; Ports: 2
- Protocol: RS-232 115 Kbps; Ports: 4
- Protocol: USB 600 Mbps; Ports: 2

Storage networking supported

- CIFS
- FC
- iSCSI
- NFS v3
- NFS v4.0
- NFS v4.1
- NFS v4.2
- NVMe/FC
- NVMe/TCP
- S3
- S3 with NAS
- SMB 2.0
- SMB 2.1
- SMB 2.x
- SMB 3.0
- SMB 3.1
- SMB 3.1.1

System environment specifications

- Typical Power: 3392 BTU/hr
- Worst-case Power: 5136 BTU/hr
- Weight: 59.5 lb, 27.0 kg
- Height: 2U
- Width: 19" IEC rack-compliant (17.7" 44.9 cm)
- Depth: 30.0" (35.2" with cable management bracket)
- Operating Temp/Altitude/Humidity: 10°C to 35°C (50°F to 95°F) at up to 3048m (10000 ft) elevation; 8% to 80% relative humidity, noncondensing
- Non-operating Temp/Humidity: -40°C to 70°C (-40°F to 158°F) up to 12192m (40000 ft); 10% to 95% relative humidity, noncondensing, in original container

- Acoustic Noise: Declared sound power (LwAd): 8.5; Sound pressure (LpAm) (bystander positions): 67.7 dB

Compliance

- Certifications EMC/EMI: AMCA, FCC, ICES, KC, Morocco, VCCI
- Certifications safety: BIS, CB, CSA, G_K_U-SoR, IRAM, NOM, NRCS, SONCAP, TBS
- Certifications Safety/EMC/EMI: EAC, UKRSEPRO
- Certifications Safety/EMC/EMI/RoHS: BSMI, CE DoC, UKCA DoC
- Standards EMC/EMI: BS-EN-55032, BS-EN55035, CISPR 32, EN55022, EN55024, EN55032, EN55035, EN61000-3-2, EN61000-3-3, FCC Part 15 Class A, ICES-003, KS C 9832, KS C 9835
- Standards Safety: ANSI/UL60950-1, ANSI/UL62368-1, BS-EN62368-1, CAN/CSA C22.2 No. 60950-1, CAN/CSA C22.2 No. 62368-1, CNS 15598-1, EN60825-1, EN62368-1, IEC 62368-1, IEC60950-1, IS 13252(part 1)

High availability

- Ethernet based baseboard management controller (BMC) and ONTAP management interface
- Redundant hot-swappable controllers
- Redundant hot-swappable power supplies
- SAS in-band management over SAS connections for external shelves

Key specifications for FAS90

The following are select specifications for the FAS90 storage system in a single high availability pair. Visit NetApp Hardware Universe (HWU) for the complete specifications for this storage system.

FAS90 specifications at a glance

- Platform Configuration: FAS90 Dual Chassis HA Pair
- Max Raw Capacity: 15.0000 PB
- Memory: 1024.0000 GB
- Form Factor: 2U chassis with 1 HA controllers
- ONTAP Version: ONTAP: 9.16.1P2
- PCIe Expansion Slots: 18
- Minimum ONTAP Version: ONTAP 9.15.1

Scaleout maximums

- Type: NAS; HA Pairs: 12; Raw Capacity: 180.0 PB / 159.9 PiB; Max Memory: 12288 GB
- Type: SAN; HA Pairs: 6; Raw Capacity: 90.0 PB / 79.9 PiB; Max Memory: 6144 GB
- Type: HA Pair; Raw Capacity: 15.0 PB / 13.3 PiB; Max Memory: 1024.0000

I/O

Onboard I/O

No onboard I/O data.

Total I/O

- Protocol: Ethernet 200 Gbps; Ports: 20
- Protocol: Ethernet 100 Gbps; Ports: 32
- Protocol: Ethernet 25 Gbps; Ports: 48
- Protocol: Ethernet 10 Gbps; Ports: 48
- Protocol: FC 64 Gbps; Ports: 48
- Protocol: NVMe/FC 64 Gbps; Ports: 48
- Ports: 0
- Protocol: SAS 12 Gbps; Ports: 48

Management ports

- Protocol: Ethernet 1 Gbps; Ports: 2
- Protocol: RS-232 115 Kbps; Ports: 4
- Protocol: USB 600 Mbps; Ports: 2

Storage networking supported

- CIFS
- FC
- iSCSI
- NFS v3
- NFS v4.0
- NFS v4.1
- NFS v4.2
- NVMe/FC
- NVMe/TCP
- S3
- S3 with NAS
- SMB 2.0
- SMB 2.1
- SMB 2.x
- SMB 3.0
- SMB 3.1
- SMB 3.1.1

System environment specifications

- Typical Power: 3392 BTU/hr
- Worst-case Power: 5136 BTU/hr
- Weight: 59.5 lb, 27.0 kg
- Height: 2U
- Width: 19" IEC rack-compliant (17.7" 44.9 cm)
- Depth: 30.0" (35.2" with cable management bracket)
- Operating Temp/Altitude/Humidity: 10°C to 35°C (50°F to 95°F) at up to 3048m (10000 ft) elevation; 8% to 80% relative humidity, noncondensing
- Non-operating Temp/Humidity: -40°C to 70°C (-40°F to 158°F) up to 12192m (40000 ft); 10% to 95% relative humidity, noncondensing, in original container
- Acoustic Noise: Declared sound power (LwAd): 8.5; Sound pressure (LpAm) (bystander positions): 67.7 dB

Compliance

- Certifications EMC/EMI: AMCA, FCC, ICES, KC, Morocco, VCCI
- Certifications safety: BIS, CB, CSA, G_K_U-SoR, IRAM, NOM, NRCS, SONCAP, TBS
- Certifications Safety/EMC/EMI: EAC, UKRSEPRO
- Certifications Safety/EMC/EMI/RoHS: BSMI, CE DoC, UKCA DoC
- Standards EMC/EMI: BS-EN-55032, BS-EN55035, CISPR 32, EN55022, EN55024, EN55032, EN55035, EN61000-3-2, EN61000-3-3, FCC Part 15 Class A, ICES-003, KS C 9832, KS C 9835
- Standards Safety: ANSI/UL60950-1, ANSI/UL62368-1, BS-EN62368-1, CAN/CSA C22.2 No. 60950-1, CAN/CSA C22.2 No. 62368-1, CNS 15598-1, EN60825-1, EN62368-1, IEC 62368-1, IEC60950-1, IS 13252(part 1)

High availability

- Ethernet based baseboard management controller (BMC) and ONTAP management interface
- Redundant hot-swappable controllers
- Redundant hot-swappable power supplies
- SAS in-band management over SAS connections for external shelves

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—with prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.