



# **AFX systems**

Install and maintain

NetApp

February 06, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap-systems/afx/afx-install-setup.html> on February 06, 2026. Always check [docs.netapp.com](https://docs.netapp.com) for the latest.

# Table of Contents

AFX system .....	1
Install and setup your AFX storage system .....	1
Maintain your AFX storage system .....	1
Overview of the maintenance procedures - AFX 1K .....	1
Boot media .....	2
Chassis .....	12
Controller .....	21
Replace a DIMM - AFX 1K .....	32
Replace a fan - AFX 1K .....	38
Replace NVRAM - AFX 1K .....	39
Replace the NV battery - AFX 1K .....	44
I/O module .....	49
Replace a power supply - AFX 1K .....	65
Replace the real-time clock battery - AFX 1K .....	67
Replace system management module - AFX 1K .....	72

# AFX system

## Install and setup your AFX storage system

Go to the [AFX documentation](#) to learn how to install, use, and manage your AFX system.

The [AFX documentation](#) includes information such as:

- Installation and setup instructions
- Administration instructions for configuring your system, such as provisioning NAS storage, cloning data, and re-sizing local storage.
- Instructions for managing your system, including managing client access, securing your data, and protecting your data.
- Monitoring and troubleshooting instructions, including information about alerts, cluster events, and system logs.

Specific maintenance procedures for each type of AFX system are in the [AFX system maintenance section](#).

## Maintain your AFX storage system

### Overview of the maintenance procedures - AFX 1K

Maintain the hardware of your AFX 1K 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 AFX 1K storage system has already been deployed as a storage node in the ONTAP environment.

#### System components

For the AFX 1K storage system, you can perform maintenance procedures on the following components.

<a href="#">Boot media</a>	The boot media stores a primary and secondary set of ONTAP image files that the system uses when it boots.
<a href="#">Controller</a>	A controller consists of a board, firmware, and software. It controls the drives and runs the ONTAP operating system software.
<a href="#">DIMM</a>	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
<a href="#">Fan</a>	A fan cools the controller.

<a href="#">NVRAM module</a>	The NVRAM (Non-Volatile Random Access Memory) module lets the controller protect and save in-flight data if the system loses power. The system ID is stored in the NVRAM module. When replaced, the controller uses the new system ID from the replacement NVRAM module.
<a href="#">NV battery</a>	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.
<a href="#">I/O module</a>	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.
<a href="#">Power supply</a>	A power supply provides a redundant power source in a controller.
<a href="#">Real-time clock battery</a>	A real-time clock battery preserves system date and time information if the power is off.
<a href="#">System management module</a>	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

### Boot media replacement workflow - AFX 1K

Get started with replacing the boot media in your AFX 1K storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

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 to replace the boot media - AFX 1K

Before replacing the boot media in your AFX 1K storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that there are no defective cluster ports on the controller, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

Before replacing the boot media, make sure to review the following requirements.

- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct controller:
  - The *impaired* controller is the controller on which you are performing maintenance.
  - The *healthy* controller is the HA partner of the impaired controller.
- There must be no faulty cluster ports on the impaired controller.

### What's next?

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

### Shut down the controller to replace the boot media - AFX 1K

Shut down the impaired controller in your AFX 1K storage system to prevent data loss and ensure system stability when replacing the boot media.

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

### About this task

- If you have a cluster with more than four nodes, it must be in quorum. To view cluster information about your nodes, use the `cluster show` command. For more information about the `cluster show` command, see [View node-level details in an ONTAP cluster](#).
- If the cluster is not in quorum or if the health or eligibility of any controller (other than the impaired controller) shows as false, 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 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 you are running ONTAP version 9.17.1 and 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 prevents 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	Take over or halt the impaired controller from the healthy controller: <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> The <i>-halt true</i> parameter brings the impaired node to the LOADER prompt.

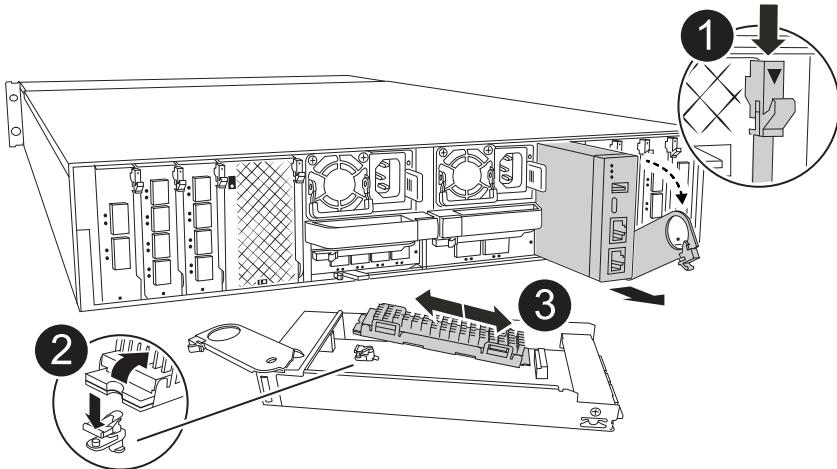
#### What's next?

After shutting down the controller, [change the boot media](#).

#### Replace the boot media - AFX 1K

The boot media in your AFX 1K 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 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.
3. Remove the System Management module:
  - a. Remove cables from the System Management module and label them to ensure correct reconnection during reinstallation.
  - 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.
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 replacing the boot media, [restore the ONTAP image from the partner node](#).

## Boot the recovery image - AFX 1K

After installing the new boot media device in your AFX 1K storage system, you can start the automated boot media recovery process to restore the configuration from the partner node.

### About this task

During the recovery process, the system checks whether encryption is enabled and identifies the type of key encryption being used. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

### Before you begin

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

### Steps

1. From the LOADER prompt, enter the command:

```
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 encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:



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

### Show example of configuration error finding prompts

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

Has key manager been configured on this system

Is the key manager onboard

If you see this message...	Do this...
key manager is not configured. Exiting.	<p>Encryption is not configured on the system. Complete the following steps:</p> <ol style="list-style-type: none"><li>Press &lt;enter&gt; when console messages stop.<ul style="list-style-type: none"><li>If you see the login prompt, go to step 4.</li><li>If you do not see login prompt, log into the partner node and proceed to step 4.</li></ul></li><li>Go to step 6 to enable automatic giveback if it was disabled.</li></ol>
key manager is configured.	<p>Go to step 5 to restore the appropriate key manager.</p> <p>The node accesses the boot menu and runs:</p> <ul style="list-style-type: none"><li>Option 10 for systems with Onboard Key Manager (OKM).</li><li>Option 11 for systems with External Key Manager (EKM).</li></ul>

4. If encryption is not installed on the system and the login prompt is not displayed. Complete the following steps:

- Give back only the root with override-destination-checks option:

```
storage failover giveback -ofnode impaired-node -only-root true -override
-destination-checks true
```



This command is available only in Diagnostic mode. For details, see [Privilege levels for ONTAP CLI commands](#).

If you encounter errors, contact [NetApp Support](#).

- Wait 5 minutes after the giveback report completes, and check failover status and giveback status:

```
storage failover show and storage failover show-giveback
```



The following command is only available in the Diagnostic mode privilege level.

- c. If you are running ONTAP 9.17.1 and the HA internconnect links were taken down, bring them back up:

```
system ha interconnect link on -node healthy-node -link 0
```

```
system ha interconnect link on -node healthy-node -link 1
```



If you are running 9.18.1 or higher, skip the above step and go to the next step.

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

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

- 5. For systems with key-manager configured, select the appropriate key manager restoration process.

## Onboard Key Manager (OKM)

If OKM is detected, 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 following when prompted:
  - i. The passphrase
  - ii. The passphrase again when prompted to confirm
  - iii. Backup data for onboard key manager

### 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 ACKUP-----
```

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

When the recovery process is complete, the node will reboot. 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.
```

- d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.
- e. Return the impaired controller to normal operation by giving back its storage:

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

- i. If the HA interconnect links were taken down, bring them back up to resume automatic giveback:

```
system ha interconnect link on -node healthy-node -link 0
```

```
system ha interconnect link on -node healthy-node -link 1
```

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

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

### External Key Manager (EKM)

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

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

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

- 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 attempts to restore the appropriate files from the partner node and reboots the node. Continue to the next step.

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

```
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>
```

- b. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.
- c. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
i. If the HA interconnect links were taken down, bring them back up to resume automatic
giveback:
```

```
system ha interconnect link on -node healthy-node -link 0
```

```
system ha interconnect link on -node healthy-node -link 1
```

6. If automatic giveback was disabled, reenable it:

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

7. 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 need to [return the failed part](#)

to NetApp.

## Return the failed part to NetApp - AFX 1K

If a component in your AFX 1K storage system fails, return the failed part to NetApp. See the [Part Return and Replacements](#) page for further information.

## Chassis

### Chassis replacement workflow - AFX 1K

Get started with replacing the chassis of your AFX 1K 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 the chassis replacement](#)

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 - AFX 1K

Before replacing the chassis in your AFX 1K storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have 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 reviewing the requirements, [get ready to replace the chassis](#).

## Prepare to replace the chassis - AFX 1K

Prepare to replace the impaired chassis in your AFX 1K storage 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 AFX 1K chassis hardware, you need to [shut down the controller](#).

### Shut down the controller to replace the chassis - AFX 1K

Shut down the controller in your AFX 1K 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, perform a storage failover takeover of the controller so that the healthy controller continues to serve data from the impaired controller storage.

#### About this task

- If you have a cluster with more than four nodes, it must be in quorum. To view cluster information about your nodes, use the `cluster show` command. For more information about the `cluster show` command, see [View node-level details in an ONTAP cluster](#).
- If the cluster is not in quorum or if the health or eligibility of any controller (other than the impaired controller) shows as false, 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 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 you are running ONTAP version 9.17.1 and 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 prevents 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	<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 the impaired node to the LOADER prompt.</p>

### What's next?

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

## Replace the chassis - AFX 1K

Replace the chassis of your AFX 1K storage system when a hardware failure requires it. The replacement process involves removing the controller, I/O cards, NVRAM12 module, system management module, and power supply units (PSUs), installing the replacement chassis, and reinstalling the chassis components.

### Step 1: Remove the PSUs and cables

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

#### Steps

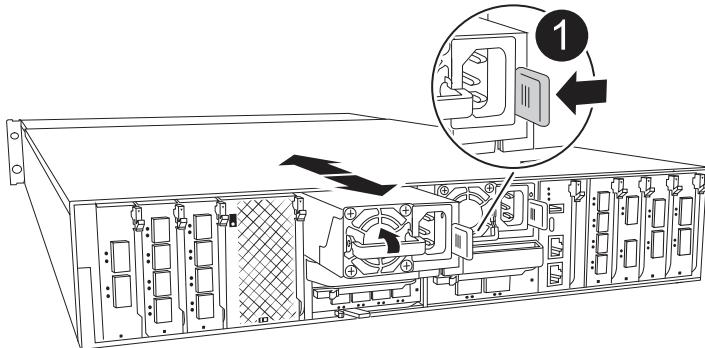
1. Remove the PSUs:
  - a. If you are not already grounded, properly ground yourself.
  - b. Unplug power cords from the PSUs.

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

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



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



1

Terracotta PSU locking tab

d. Repeat these steps for the second PSU.

2. Remove the cables:

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

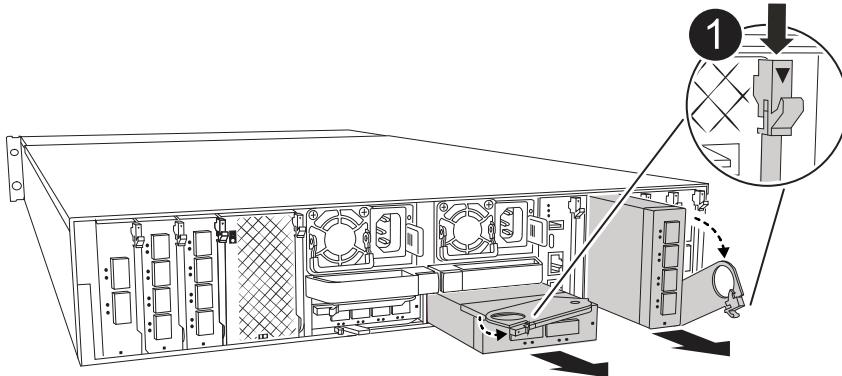


Cables should have been labeled at the beginning of this procedure.

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

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

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



1

I/O cam latch

a. Depress the cam button on the target module.

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

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

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

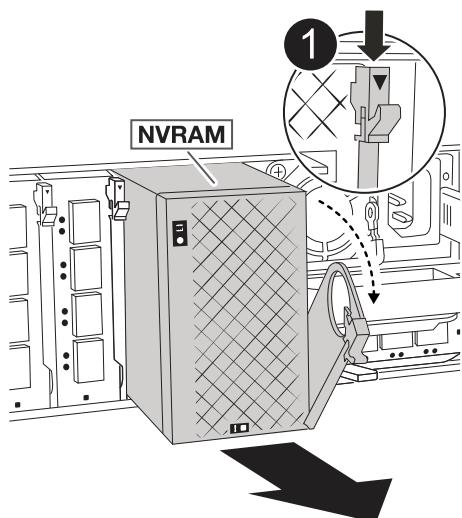
- d. Set the I/O module aside and repeat these steps for any other I/O modules.

2. Remove the NVRAM12 module:

- a. Depress the locking cam button.

The cam button moves away from the chassis.

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

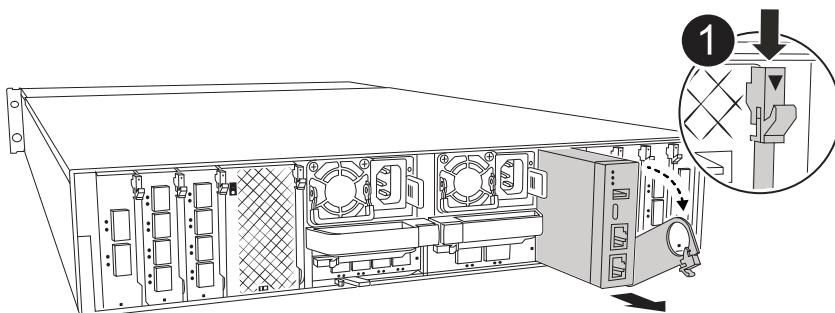


1	NVRAM12 cam latch
---	-------------------

- d. Set the NVRAM module on a stable surface.

3. Remove the system management module:

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

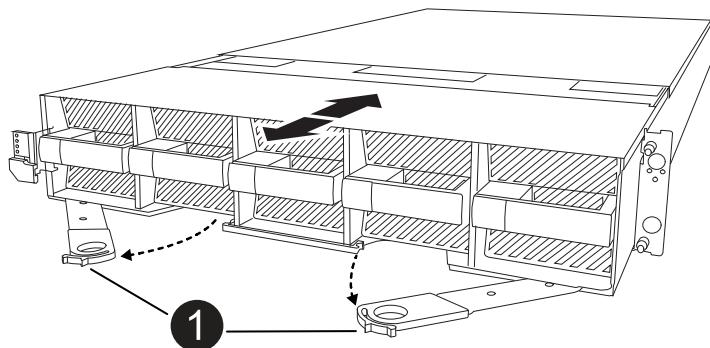


1	System Management module cam latch
---	------------------------------------

### Step 3: Remove the controller module

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

The controller module moves slightly out of the chassis.



<b>1</b>	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 replacing the impaired AFF A1K chassis and reinstalling its components, complete the [chassis replacement](#).

#### Complete the chassis replacement - AFX 1K

Reboot the controller, verify system health, and return the failed part to NetApp to complete the final step in the AFX 1K chassis replacement procedure.

## Step 1: Boot the controllers and give back the controllers

After the controllers reboot, boot ONTAP and give back the controllers.

### Steps

1. Check the console output:
  - a. If the controller stop at the LOADER prompt, boot the controller with the `boot_ontap` command.
  - b. Once the node boots and provides the login prompt, log into the partner controller and check that the replaced controller is ready for giveback with `storage failover show` command.
2. Perform the giveback:
  - a. Connect the console cable to the partner controller.
  - b. Press `<enter>` when console messages stop.
    - If you see the `login` prompt, go to the next step.
    - If you do not see login prompt, log into the partner node.
  - c. Give back only the root with `override-destination-checks` option:

```
storage failover giveback -ofnode impaired-node -only-root true -override-destination-checks true
```



The following command is only available in the Diagnostic mode privilege level. For more information on privilege levels, see [Understand the privilege levels for ONTAP CLI commands](#).

If you encounter errors, contact [NetApp Support](#).

- d. Wait 5 minutes after the giveback report completes, and check failover status and giveback status:

```
storage failover show and storage failover show-giveback
```



The following command is only available in the Diagnostic mode privilege level.

- e. If the HA internconnect links were taken down, bring them back up:

```
system ha interconnect link on -node healthy-node -link 0
```

```
system ha interconnect link on -node healthy-node -link 1
```

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

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

- g. If automatic giveback was disabled, reenable it: `storage failover modify -node impaired_node_name -auto-giveback-of true`
- h. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

## Step 2: Verify storage system health

After the controller gives back the storage, check the health using [Active IQ Config Advisor](#).

### Steps

1. After the giveback is complete, run Active IQ Config Advisor to verify the health of the storage system.
2. Correct any issues you encounter.

## Step 3: Return the failed part to NetApp

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

## Controller

### Controller replacement workflow - AFX 1K

Get started with replacing the controller in your AFX 1K storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and verifying system operations.

1

#### [Review the requirements to replace the controller](#)

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](#)

Replace the controller by removing the impaired controller, moving the FRU components to the replacement module, and installing the replacement module in the enclosure.

4

#### [Restore and verify the system configuration](#)

Verify the low-level system configuration of the replacement controller and update the system settings if needed.

5

#### [Give back the controller](#)

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

6

#### [Complete controller replacement](#)

Verify the Logical Interfaces (LIFs), check cluster health, and return the failed part to NetApp.

## Requirements to replace the controller - AFX 1K

Before replacing the controller in your AFX 1K storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

Review the requirements for replacing the controller.

- 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.
- 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").
- 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.
- 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 reviewing the requirements to replace your AFX 1K controller, [turn off the controllers](#).

## Shut down the impaired controller - AFX 1K

Shut down the impaired controller in your AFX 1K storage system to prevent data loss and ensure system stability when replacing the controller.

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

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

### About this task

- If you have a cluster with more than four nodes, it must be in quorum. To view cluster information about your nodes, use the `cluster show` command. For more information about the `cluster show` command, see [View node-level details in an ONTAP cluster](#).
- If the cluster is not in quorum or if the health or eligibility of any controller (other than the impaired controller) shows as false, 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 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 you are running ONTAP version 9.17.1 and 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 prevents 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	Take over or halt the impaired controller from the healthy controller: <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> The <b>-halt true</b> parameter brings the impaired node to the LOADER prompt.

## What's next?

After shutting down the controller, [replace the controller](#).

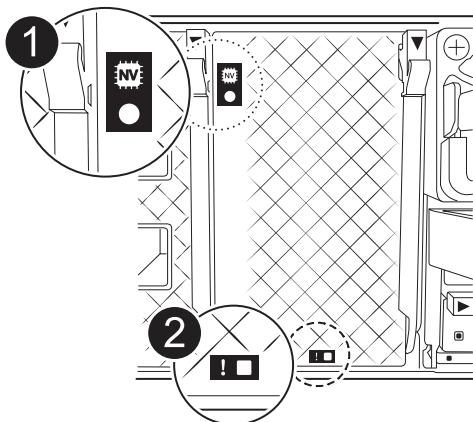
## Replace the controller - AFX 1K

Replace the controller in your AFX 1K storage system when a hardware failure requires it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

### 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.

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:

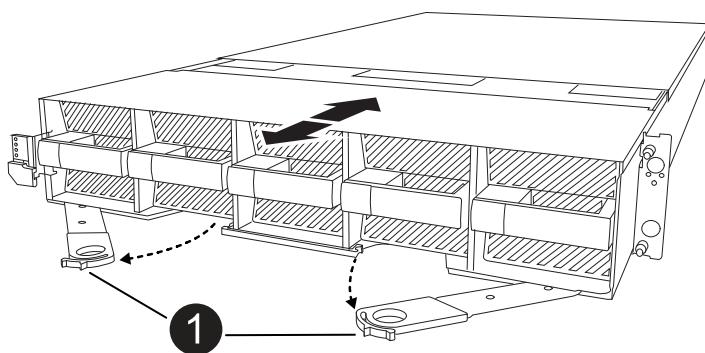


1	NVRAM status LED
2	NVRAM attention LED

- 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. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel and pulling towards you until the bezel releases from the ball studs on the chassis frame.
4. 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
---	---------------------

5. 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.

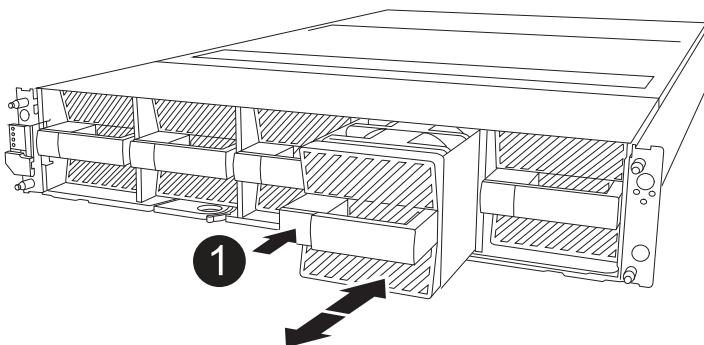
## Step 2: Move the fans

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

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



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



1

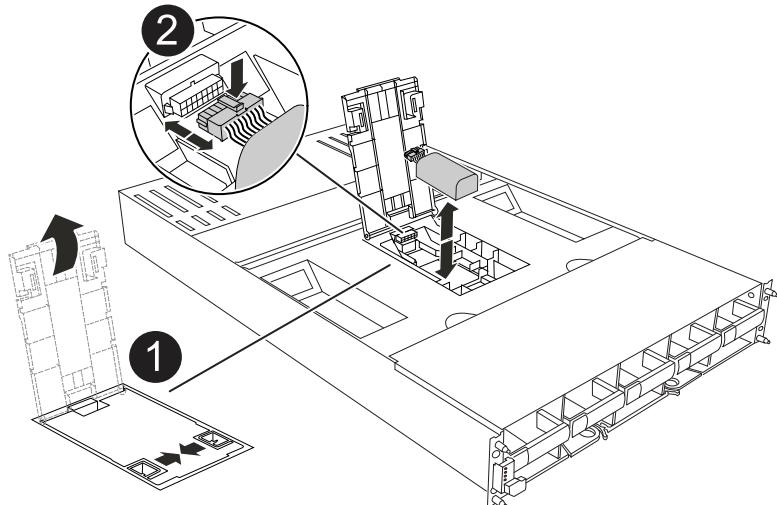
Black locking button

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

## Step 3: Move the NV battery

Move the NV battery to the replacement controller.

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



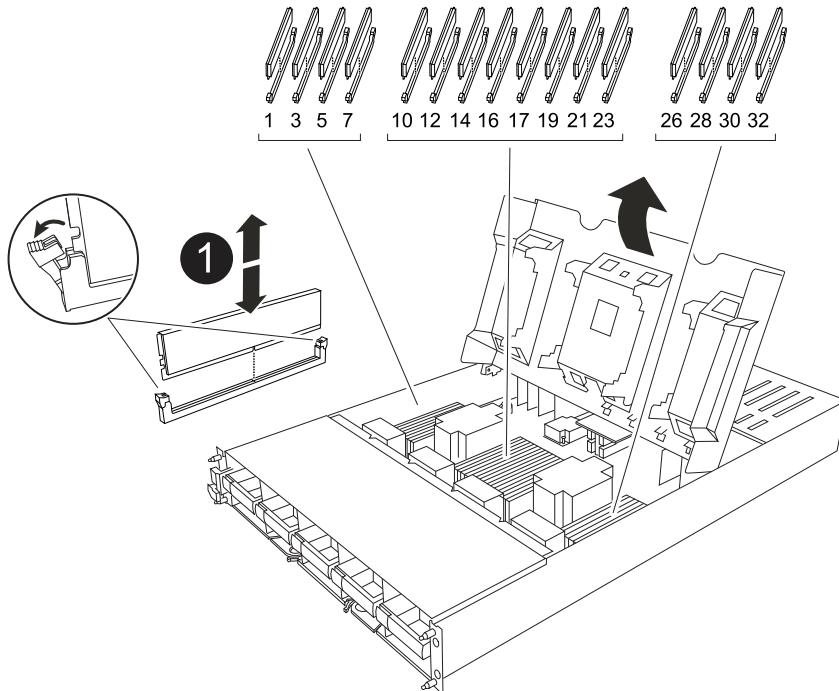
1	NV battery air duct cover
2	NV battery plug
3	NV battery pack

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

#### Step 4: Move system DIMMs

Move the DIMMs to the replacement controller module.

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 you should be able to insert it 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 is displayed in local time and 24-hour format.

5. If necessary, set the current date with the `set date mm/dd/yyyy` command.
6. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
  - a. You can get the current GMT from the partner node with the `date -u` command.

## What's next?

After replacing the impaired AFX 1K controller, [restore the system settings](#).

## Restore and verify the system configuration - AFX 1K

Verify that the controller's HA configuration is active and functioning correctly in your AFX 1K 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.

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.

### Step 2: Verify disk list

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

```
storage show disk -p
```

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

## 2. Exit Maintenance mode:

```
halt
```

### What's next?

After restoring and verifying the AFX 1K storage system configuration, [give back the controller](#).

### Give back the controller - AFX 1K

Return control of storage resources to the replacement controller so your AFX 1K storage 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 do not see login prompt, press the <enter> key, if still no prompt, log into the partner node.

3. Give back only the root with override-destination-checks option:

```
storage failover giveback -ofnode impaired-node -only-root true -override-destination-checks true
```



The following command is only available in the Diagnostic mode privilege level. For more information on privilege levels, see [Understand the privilege levels for ONTAP CLI commands](#).

If you encounter errors, contact [NetApp Support](#).

4. Wait 5 minutes after the giveback report completes, then check failover and giveback status:

```
storage failover show and storage failover show-giveback
```



The following command is only available in the Diagnostic mode privilege level.

5. If the HA internconnect links were taken down, bring them back up:

```
system ha interconnect link on -node healthy-node -link 0
```

```
system ha interconnect link on -node healthy-node -link 1
```

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

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

7. If automatic giveback was disabled, reenable it:

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

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

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

## OKM encryption

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 can get this passphrase from the customer, or contact [NetApp Support](#).



You will be 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. Move the console cable to the partner node and enter the following login:

```
admin
```

7. Give back only the root with `override-destination-checks` option:

```
storage failover giveback -ofnode impaired-node -only-root true -override-destination-checks true
```



The following command is only available in the Diagnostic mode privilege level. For more information on privilege levels, see [Understand the privilege levels for ONTAP CLI commands](#).

If you encounter errors, contact [NetApp Support](#).

8. Wait 5 minutes after the giveback report completes, then check failover and giveback status:

```
storage failover show and storage failover show-giveback
```



The following command is only available in the Diagnostic mode privilege level.

9. Move the console cable to the replacement node and enter the following:

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



You will be prompted for the cluster-wide passphrase of OKM for the cluster.

10. Check status of the keys with the following command:

```
security key-manager key query -key-type svm-KEK
```

If the *Restored* column shows anything but *true*, contact [NetApp Support](#).

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

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

- a. If the HA interconnect links were taken down, bring them back up to resume automatic giveback:

```
system ha interconnect link on -node healthy-node -link 0
```

```
system ha interconnect link on -node healthy-node -link 1
```

12. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback-of 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 transferring storage resource ownership to the replacement controller, [finish replacing the controller](#).

## Complete controller replacement - AFX 1K

Verify that the logical interfaces (LIFs) are reporting to their home port, perform a cluster health check, and then return the failed part to NetApp to complete the final step in the AFX 1K controller replacement procedure.

### Step 1: Verify LIFs and check cluster health

Before returning the replacement node to service, ensure the logical interfaces are on their home ports, check 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, return 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 for more information.

### 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 - AFX 1K

Replace a DIMM in your AFX 1K storage system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

## Before you begin

- Make sure you have the replacement component you received from NetApp.
- 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.

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

### About this task

- If you have a cluster with more than four nodes, it must be in quorum. To view cluster information about your nodes, use the `cluster show` command. For more information about the `cluster show` command, see [View node-level details in an ONTAP cluster](#).
- If the cluster is not in quorum or if the health or eligibility of any controller (other than the impaired controller) shows as false, 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 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 you are running ONTAP version 9.17.1 and 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 prevents 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:

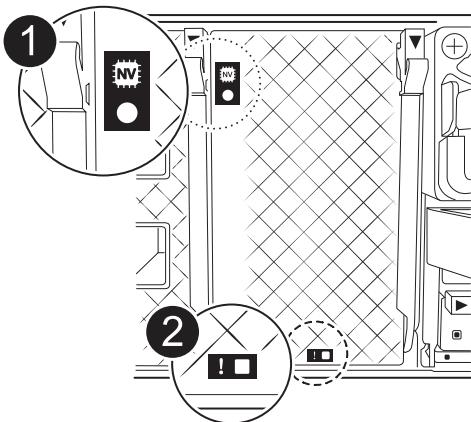
If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.

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 the impaired node 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.

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:

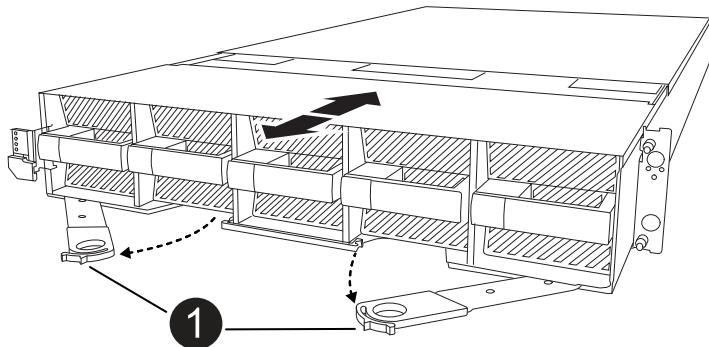


1	NVRAM status LED
2	NVRAM attention LED

- 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. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel and pulling towards you until the bezel releases from the ball studs on the chassis frame.
4. 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

5. 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.

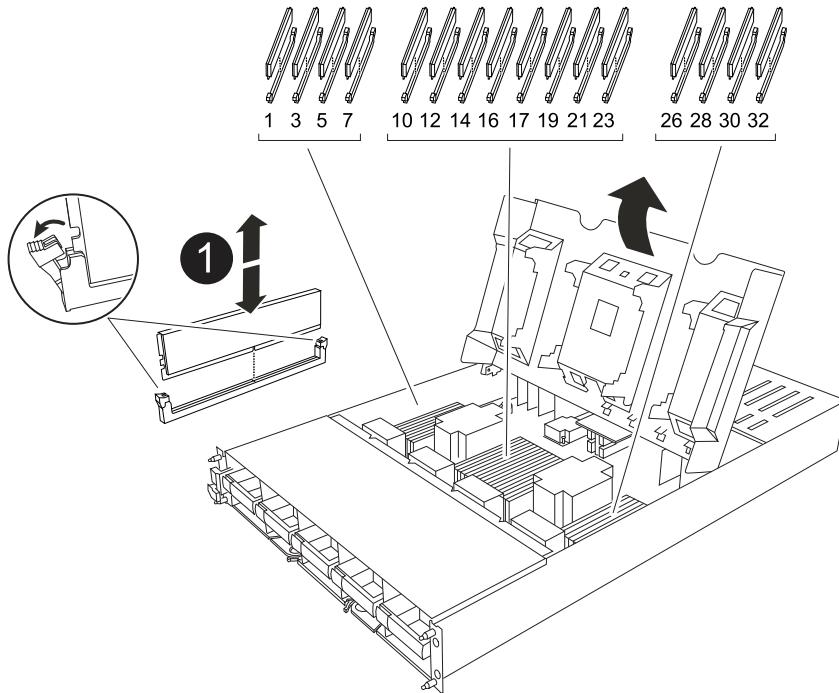
1. If you are not already grounded, properly ground yourself.
2. Open the controller air duct on the top of the controller.
  - a. Insert your fingers in the recesses at the far ends of the air duct.
  - b. Lift the air duct and rotate it upward as far as it will go.
3. Locate the DIMMs on your controller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

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

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



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.

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. Align the bezel with the ball studs and then gently push the bezel into place.
5. Press <enter> when console messages stop.
  - If you see the login prompt, go to the next step.
  - If you do not see a login prompt, log into the partner node.
6. Give back only the root with override-destination-checks option:

```
storage failover giveback -ofnode impaired-node -only-root true -override-destination-checks true
```



The following command is only available in the Diagnostic mode privilege level. For more information on privilege levels, see [Understand the privilege levels for ONTAP CLI commands](#).

If you encounter errors, contact [NetApp Support](#).

7. Wait 5 minutes after the giveback report completes, then check failover and giveback status:

```
storage failover show and storage failover show-giveback
```



The following command is only available in the Diagnostic mode privilege level.

8. If the HA internconnect links were taken down, bring them back up:

```
system ha interconnect link on -node healthy-node -link 0
```

```
system ha interconnect link on -node healthy-node -link 1
```

9. If automatic giveback was disabled, reenable it:

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

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

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

11. 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 - AFX 1K

Replace a failed or faulty fan module in your AFX 1K system to maintain proper cooling and prevent system performance issues. The fans are hot-swappable and can be replaced without shutting down the system. This procedure includes identifying the faulty fan using console error messages and LED indicators, removing the bezel, swapping the fan module, 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.

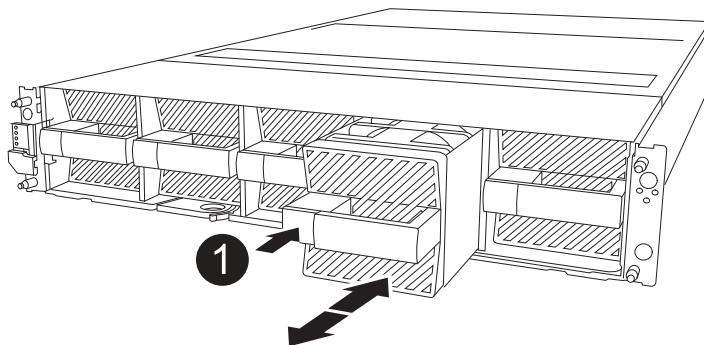


There is a single LED for each fan. It is green when the fan is functioning correctly and amber when not.

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.



	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 NVRAM - AFX 1K

Replace the NVRAM in your AFX 1K storage 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, 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.

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

### About this task

- If you have a cluster with more than four nodes, it must be in quorum. To view cluster information about your nodes, use the `cluster show` command. For more information about the `cluster show` command, see [View node-level details in an ONTAP cluster](#).
- If the cluster is not in quorum or if the health or eligibility of any controller (other than the impaired controller) shows as false, 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 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 you are running ONTAP version 9.17.1 and 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 prevents 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	<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 the impaired node 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.

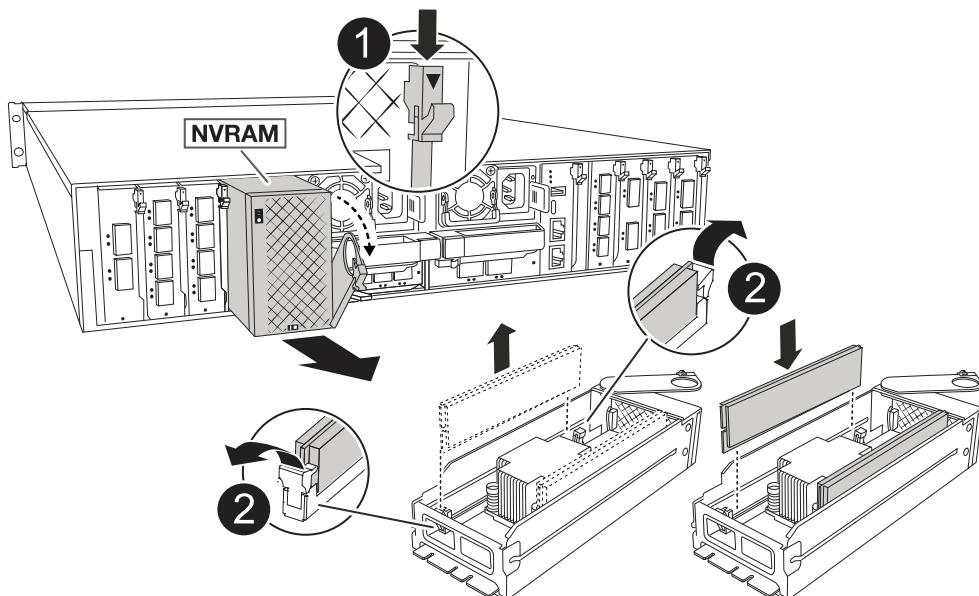
### Option 1: Replace the NVRAM module

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

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

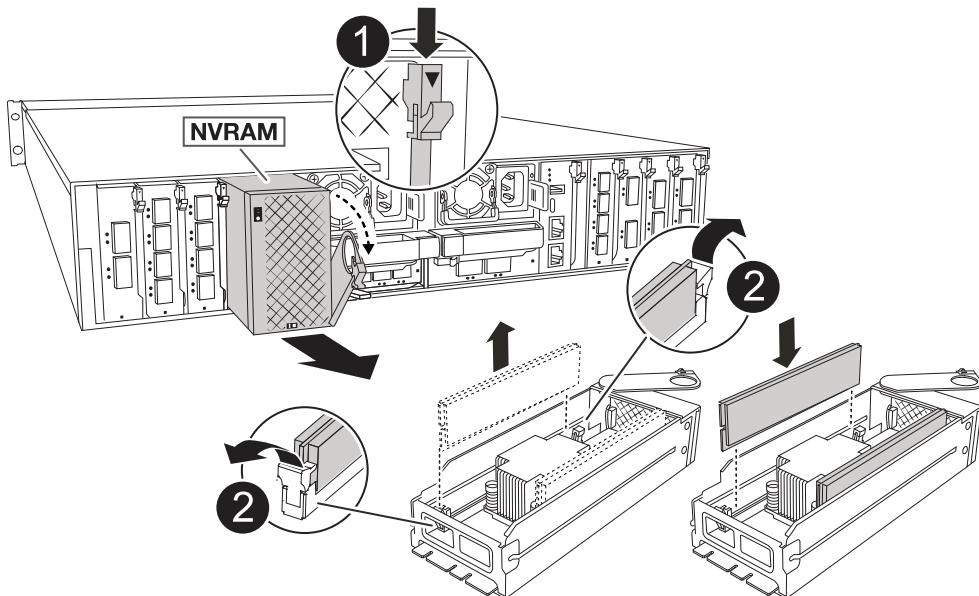
5. Set the NVRAM module on a stable surface.
6. Remove the DIMMs, one at a time, from the impaired NVRAM module and install them in the replacement NVRAM module.
7. Install the replacement NVRAM module into the enclosure:
  - a. Align the module with the edges of the enclosure opening in slot 4/5.
 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

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. 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.

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.

### Step 4: Complete NVRAM replacement

Perform the following steps to complete the NVRAM replacement.

#### Steps

1. 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.

The following is an example of the prompt for overriding the system ID:

```
WARNING: System ID mismatch. This usually occurs when replacing a boot
device or NVRAM cards!
Override system ID? {y|n}
```

2. 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 displaying the current state of the storage replacement. In the following example, node2 has undergone replacement and displays the current state as "In takeover".

```
node1:> storage failover show
                                         Takeover
      Node          Partner          Possible      State Description
      -----        -----          -----
      -----
      node1        node2        false      In takeover
      node2        node1          -      Waiting for giveback
```

3. Give back the controller:

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

The controller takes back its storage and completes booting.



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

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

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

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

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

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

5. Press <enter> when console messages stop.

- If you see the *login* prompt, go to the next step.
- If you do not see login prompt, log into the partner node.

6. Wait 5 minutes after the giveback report completes, and check failover status and giveback status:

```
storage failover show and storage failover show-giveback
```



The following command is only available in the Diagnostic mode privilege level.

7. If the HA internconnect links were taken down, bring them back up:

```
system ha interconnect link on -node healthy-node -link 0
```

```
system ha interconnect link on -node healthy-node -link 1
```

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

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

9. If automatic giveback was disabled, reenable it:

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

10. 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 - AFX 1K

Replace the NV battery in your AFX 1K storage system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power

outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

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.

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

#### About this task

- If you have a cluster with more than four nodes, it must be in quorum. To view cluster information about your nodes, use the `cluster show` command. For more information about the `cluster show` command, see [View node-level details in an ONTAP cluster](#).
- If the cluster is not in quorum or if the health or eligibility of any controller (other than the impaired controller) shows as false, 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 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 you are running ONTAP version 9.17.1 and 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 prevents 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:

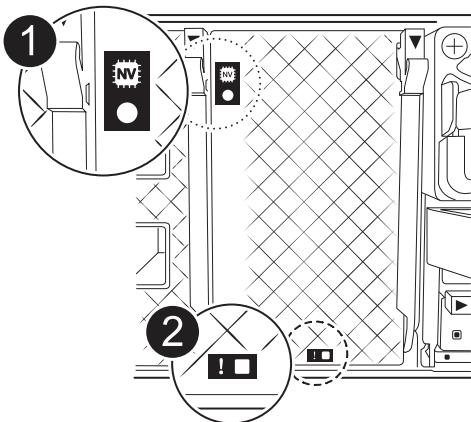
If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.

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 the impaired node 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.

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:

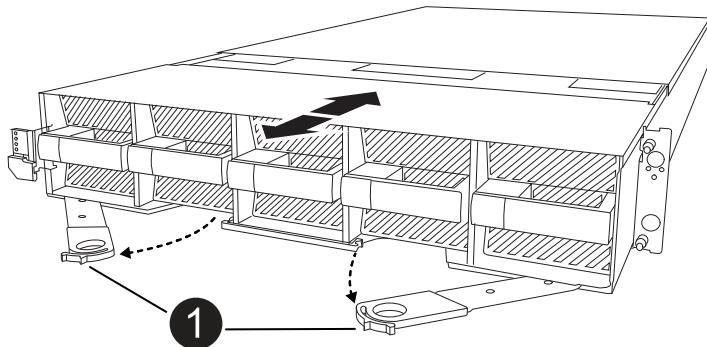


1	NVRAM status LED
2	NVRAM attention LED

- 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. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel and pulling towards you until the bezel releases from the ball studs on the chassis frame.
4. 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
---	---------------------

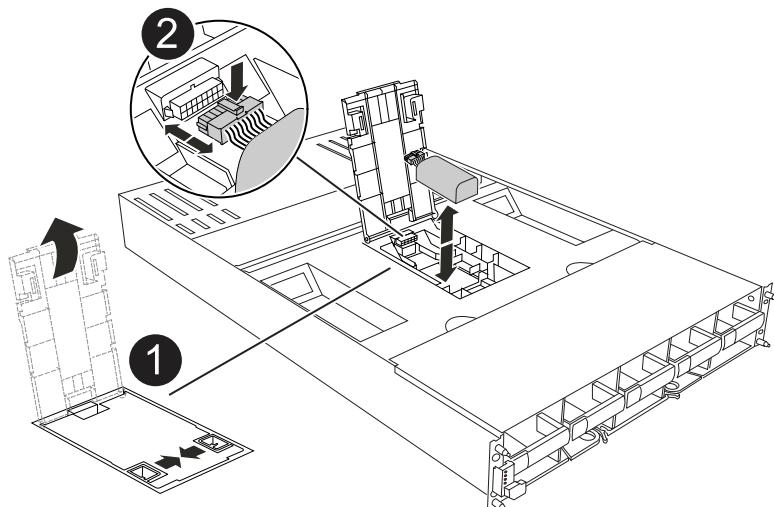
5. Slide the controller module out of the enclosure and place it on a flat, stable surface.

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.

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.

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. Align the bezel with the ball studs and then gently push the bezel into place.
5. Press <enter> when console messages stop.
  - If you see the login prompt, go to the next step.
  - If you do not see a login prompt, log into the partner node.
6. Give back only the root with override-destination-checks option:

```
storage failover giveback -ofnode impaired-node -only-root true -override-destination-checks true
```



The following command is only available in the Diagnostic mode privilege level. For more information on privilege levels, see [Understand the privilege levels for ONTAP CLI commands](#).

If you encounter errors, contact [NetApp Support](#).

7. Wait 5 minutes after the giveback report completes, then check failover and giveback status:

```
storage failover show and storage failover show-giveback
```



The following command is only available in the Diagnostic mode privilege level.

8. If the HA internconnect links were taken down, bring them back up:

```
system ha interconnect link on -node healthy-node -link 0
```

```
system ha interconnect link on -node healthy-node -link 1
```

9. If automatic giveback was disabled, reenable it:

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

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

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

11. 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 I/O module - AFX 1K

The AFX 1K storage 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 AFX 1K storage system with the same type of I/O module, or with a different kind of I/O module. You can also add an I/O module into a system with empty slots.

- [Add an I/O module](#)

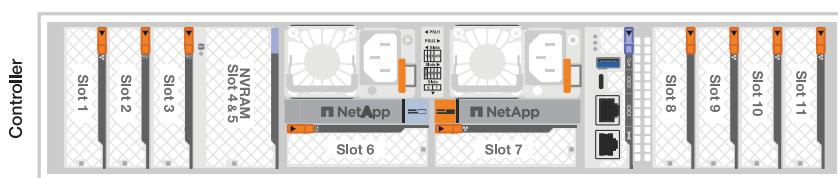
Adding additional modules can improve redundancy, helping to ensure that the system remains operational even if one module fails.

- [Replace an I/O module](#)

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

### I/O slot numbering

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



## Add I/O module - AFX 1K

Add an I/O module to your AFX 1K storage system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your AFX 1K storage system when there are empty slots available or when all slots are fully populated.

### Step 1: Shut down the impaired controller module

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

## Step 2: Add the new I/O module

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

### Before you begin

- Check the [NetApp Hardware Universe](#) to ensure the new I/O module is compatible with your storage system and ONTAP version.
- 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. Press `<enter>` when console messages stop.
  - If you see the *login* prompt, go to the next step.
  - If you do not see the login prompt, log in to the partner node.
9. Give back only the root with `override-destination-checks` option:

```
storage failover giveback -ofnode impaired-node -only-root true -override-destination-checks true
```



The following command is only available in the Diagnostic mode privilege level. For more information on privilege levels, see [Understand the privilege levels for ONTAP CLI commands](#).

If you encounter errors, contact [NetApp Support](#).

10. Wait 5 minutes after the giveback report completes, and check failover status and giveback status:

```
storage failover show and storage failover show-giveback
```



The following command is only available in the Diagnostic mode privilege level.

11. If the HA internconnect links were taken down, bring them back up:

```
system ha interconnect link on -node healthy-node -link 0
```

```
system ha interconnect link on -node healthy-node -link 1
```

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

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

13. Repeat these steps for controller B.

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

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

15. 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 <a href="#">Migrating a LIF</a> for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in <a href="#">Migrating a LIF</a> .

#### 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.

10. Press `<enter>` when console messages stop.
  - If you see the `login` prompt, go to the next step.
  - If you do not see the login prompt, log in to the partner node.
11. Give back only the root with `override-destination-checks` option:

```
storage failover giveback -ofnode impaired-node -only-root true -override-destination-checks true
```



The following command is only available in the Diagnostic mode privilege level. For more information on privilege levels, see [Understand the privilege levels for ONTAP CLI commands](#).

If you encounter errors, contact [NetApp Support](#).

12. Wait 5 minutes after the giveback report completes, and check failover status and giveback status:

```
storage failover show and storage failover show-giveback
```



The following command is only available in the Diagnostic mode privilege level.

13. If the HA internconnect links were taken down, bring them back up:

```
system ha interconnect link on -node healthy-node -link 0
system ha interconnect link on -node healthy-node -link 1
```

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

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

15. Enable automatic giveback if it was disabled:

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

16. Do one of the following:

- If you removed a storage 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 NX224 shelves, as described in [Hot-add NX224 shelf](#).

17. Repeat these steps for controller B.

### Hot swap an I/O module - AFX 1K

You can hot swap an Ethernet I/O module in your AFX 1K 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"> <li>1. Enter the following command from the console of the controller that took over its partner:</li> </ol> <pre>storage failover modify -node local -auto -giveback false</pre> <ol style="list-style-type: none"> <li>2. Enter <b>y</b> when you see the prompt <i>Do you want to disable auto-giveback?</i></li> </ol>
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: Hot swap the failed I/O module**

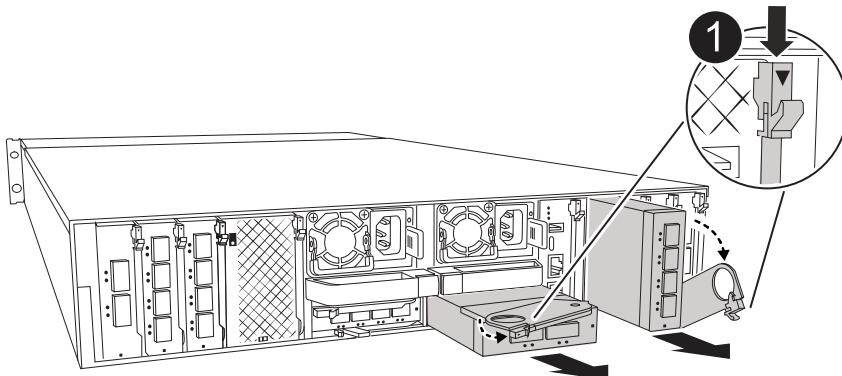
Hot swap the failed I/O module with an equivalent I/O module.

**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"> <li>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> </li> <li>2. Restore automatic giveback from the console of the controller that was taken over:  <code>storage failover modify -node local -auto -giveback true</code> </li> </ol>

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 I/O module - AFX 1K

Replace an I/O module in your AFX 1K storage system when the module fails. 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.

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

#### About this task

- If you have a cluster with more than four nodes, it must be in quorum. To view cluster information about your nodes, use the `cluster show` command. For more information about the `cluster show` command, see [View node-level details in an ONTAP cluster](#).
- If the cluster is not in quorum or if the health or eligibility of any controller (other than the impaired controller) shows as false, 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 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 you are running ONTAP version 9.17.1 and 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 prevents 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	Take over or halt the impaired controller from the healthy controller: <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> The <b>-halt true</b> parameter brings the impaired node to the LOADER prompt.

## 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.

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.

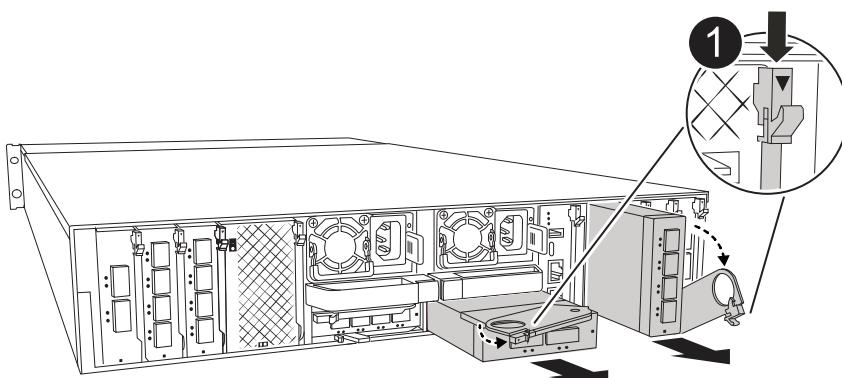


Make sure that you label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.

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. Re-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.

1. Reboot the controller from the LOADER prompt:

bye



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

2. Press <enter> when console messages stop.
  - If you see the *login* prompt, go to the next step.

- If you do not see login prompt, log into the partner node.

3. Give back only the root with override-destination-checks option:

```
storage failover giveback -ofnode impaired-node -only-root true -override
-destination-checks true
```



The following command is only available in the Diagnostic mode privilege level. For more information on privilege levels, see [Understand the privilege levels for ONTAP CLI commands](#).

If you encounter errors, contact [NetApp Support](#).

4. Wait 5 minutes after the giveback report completes, and check failover status and giveback status:

```
storage failover show and storage failover show-giveback
```



The following command is only available in the Diagnostic mode privilege level.

5. If the HA internconnect links were taken down, bring them back up:

```
system ha interconnect link on -node healthy-node -link 0
```

```
system ha interconnect link on -node healthy-node -link 1
```

6. If automatic giveback was disabled, reenable it:

```
storage failover modify -ofnode impaired-node -automatic-giveback true
```

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

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

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

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

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

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

#### Step 4: Return the failed part to NetApp

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

### Replace a power supply - AFX 1K

Replace an AC power supply unit (PSU) in your AFX 1K storage system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the target PSU, unplugging the power cable, removing the faulty PSU and installing the replacement

PSU, and then reconnecting it to the power source.

#### 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.

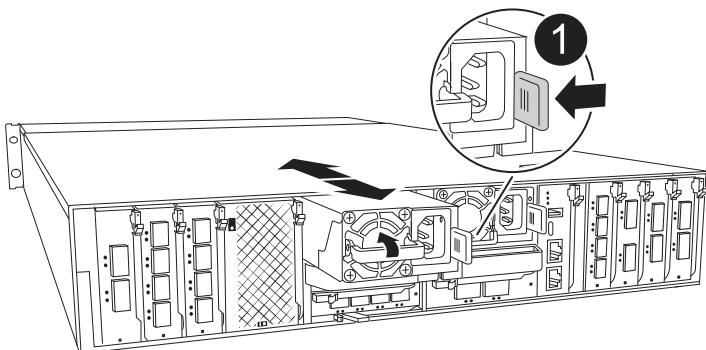
- The power supplies are redundant and hot-swappable; you do not have to takeover the controller to perform this task.

#### Steps

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



The PSU is compact. Use both hands to support it during removal to prevent it from swinging free from the controller module and causing injury.



Terracotta PSU locking tab

- Install the replacement PSU in the controller module:

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

- 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.

## Replace the real-time clock battery - AFX 1K

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your AFX 1K storage system to ensure that services and applications relying on accurate time synchronization remain operational.

### Before you begin

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

### Step 1: Shut down the impaired controller

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

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

### About this task

- If you have a cluster with more than four nodes, it must be in quorum. To view cluster information about your nodes, use the `cluster show` command. For more information about the `cluster show` command, see [View node-level details in an ONTAP cluster](#).
- If the cluster is not in quorum or if the health or eligibility of any controller (other than the impaired controller) shows as false, 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 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 you are running ONTAP version 9.17.1 and 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 prevents 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
```

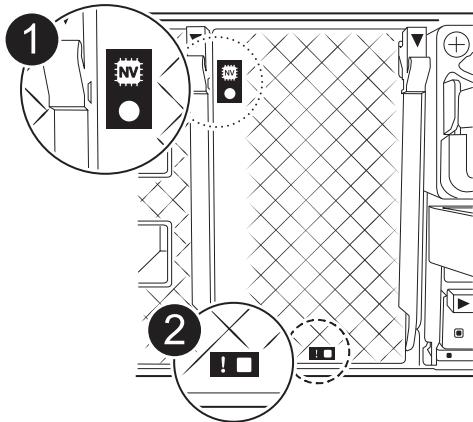
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: <code>storage failover takeover -ofnode impaired_node_name -halt true</code>  The <code>-halt true</code> parameter brings the impaired node to the LOADER prompt.

## Step 2: Remove the controller module

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

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:

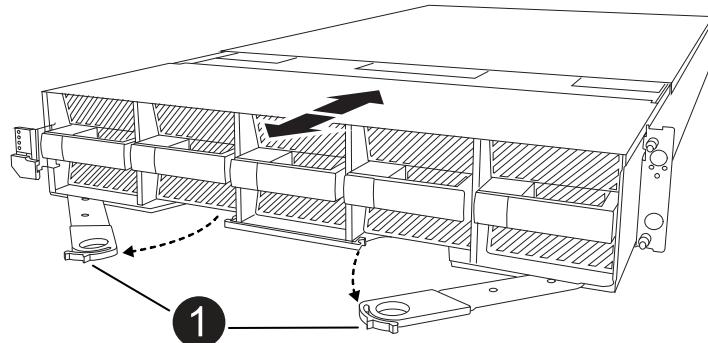


1	NVRAM status LED
2	NVRAM attention LED

- 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. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel and pulling towards you until the bezel releases from the ball studs on the chassis frame.
4. 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
---	---------------------

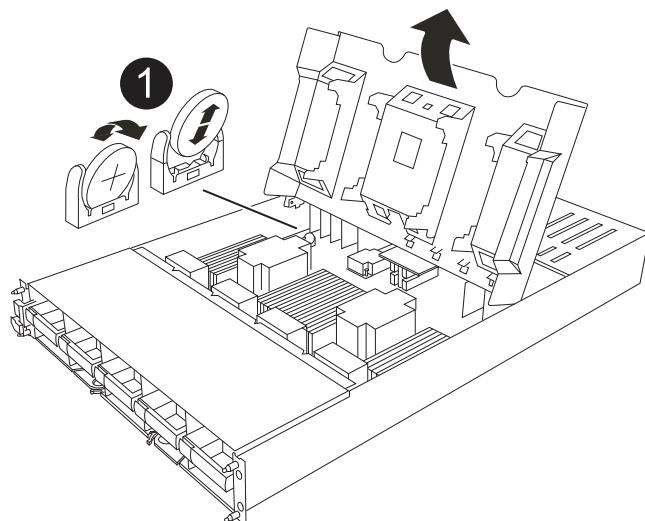
5. Slide the controller module out of the enclosure and place it on a flat, stable surface.

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 the failed RTC battery and install the replacement.

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.

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. Align the bezel with the ball studs and then gently push the bezel into place.

#### 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.

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 `show date`

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.
  - i. You can get the current GMT from the partner node with the `date -u` command.

2. Confirm the date and time on the target controller.
3. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
4. Press `<enter>` when console messages stop.
  - If you see the *login* prompt, go to the next step.
  - If you do not see login prompt, log into the partner node.
5. Give back only the root with `override-destination-checks` option:

```
storage failover giveback -ofnode impaired-node -only-root true -override-destination-checks true
```



The following command is only available in the Diagnostic mode privilege level. For more information on privilege levels, see [Understand the privilege levels for ONTAP CLI commands](#).

If you encounter errors, contact [NetApp Support](#).

6. Wait five minutes after the giveback report completes, then check the failover and giveback statuses:

```
storage failover show and storage failover show-giveback
```



The following command is only available in the Diagnostic mode privilege level.

7. If the HA internconnect links were taken down, bring them back up:

```
system ha interconnect link on -node healthy-node -link 0
```

```
system ha interconnect link on -node healthy-node -link 1
```

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

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

9. If automatic giveback was disabled, reenable it:

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

10. If AutoSupport is enabled, restore/unsuppress 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 system management module - AFX 1K

Replace the System Management module in your AFX 1K storage 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 to the left of slot 8, contains onboard components for system management, as well as ports for external management. The target controller must be shut down and powered off 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)
- RJ45 Ethernet service port for e0M and BMC network access

### 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.

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

### About this task

- If you have a cluster with more than four nodes, it must be in quorum. To view cluster information about your nodes, use the `cluster show` command. For more information about the `cluster show` command, see [View node-level details in an ONTAP cluster](#).

- If the cluster is not in quorum or if the health or eligibility of any controller (other than the impaired controller) shows as false, 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 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 you are running ONTAP version 9.17.1 and 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 prevents 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	<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 <b>-halt true</b> parameter brings the impaired node to the LOADER prompt.</p>

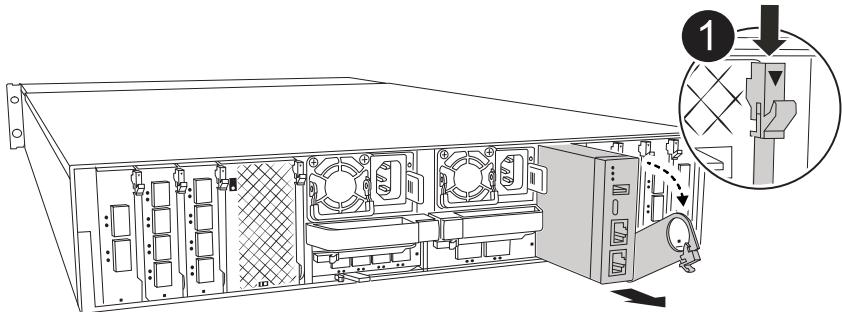
## Step 2: Replace the impaired System Management module

Replace the impaired system management module.

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.

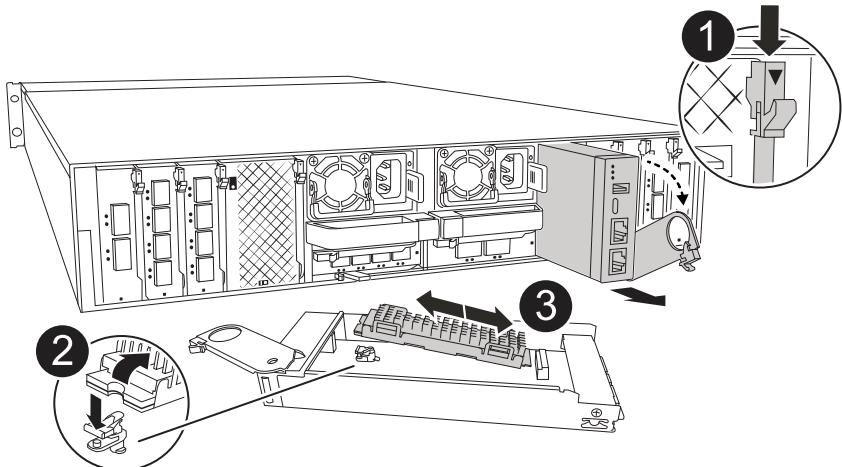


1

System Management module cam latch

- a. If you are not already grounded, properly ground yourself.
- b. Unplug the power supply cables from the PSUs.
- c. Remove any cables connected to the System Management module. Label the cables where they were connected, so you can reconnect them to the correct ports when reinstalling the module.
- d. 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.
- e. Depress the cam button on the System Management module.
- f. Rotate the cam lever down as far as it will go.
- g. Loop your finger into the hole on the cam lever and pull the module straight out of the system.
- h. Place the System Management module on an anti-static mat to access the boot media.

2. 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.

3. 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.
4. 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.
5. Rotate the cable management arm up to the closed position.
6. Recable the System Management module.

### Step 3: Reboot the controller module

Reboot the controller module.

1. Plug the power cables back into the PSU.

The system reboots, typically to the LOADER prompt.

2. Enter *bye* at the LOADER prompt.
3. Press <enter> when console messages stop.
  - If you see the *login* prompt, go to the next step.
  - If you do not see login prompt, log into the partner node.
4. Give back only the root with *override-destination-checks* option:

```
storage failover giveback -ofnode impaired-node -only-root true -override
-destination-checks true
```



The following command is only available in the Diagnostic mode privilege level. For more information on privilege levels, see [Understand the privilege levels for ONTAP CLI commands](#).

If you encounter errors, contact [NetApp Support](#).

5. Wait 5 minutes after the giveback report completes, then check the failover and giveback statuses:

```
storage failover show and storage failover show-giveback
```



The following command is only available in the Diagnostic mode privilege level.

6. If the HA internconnect links were taken down, bring them back up:

```
system ha interconnect link on -node healthy-node -link 0
```

```
system ha interconnect link on -node healthy-node -link 1
```

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

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

8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback-of true`.
9. 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

When you replace the system management module, the controller's system serial number (SSN) changes. 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

A NetApp License File (NLF) for the new system serial number is required. For more information on NetApp License Files, see [ONTAP 9.10.1 and later licensing overview](#).

You have 90 days to install the license keys, after which all old licenses become invalid. After you install a valid license key, you have 24 hours to install all of the keys before the grace period ends.

 If your system was initially running ONTAP 9.15.1 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 through ONTAP System Manager.

For more information, see [Enable new features by adding license keys with ONTAP System Manager](#).

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.

## Copyright information

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

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

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

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

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

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

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

## Trademark information

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