



# Maintain

Install and maintain

NetApp

February 06, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap-systems/asa-r2-a1k/maintain-overview.html> on February 06, 2026. Always check [docs.netapp.com](https://docs.netapp.com) for the latest.

# Table of Contents

Maintain ASA r2 systems .....	1
ASAA1K systems .....	1
Overview of the maintenance procedures - ASAA1K .....	1
Boot media .....	2
Chassis .....	13
Controller .....	21
Replace a DIMM - ASAA1K .....	32
Replace a fan - ASAA1K .....	36
Replace NVRAM - ASAA1K .....	37
Replace the NV battery - ASAA1K .....	42
I/O module .....	46
Hot-swap a power supply - ASAA1K .....	61
Replace the real-time clock battery - ASAA1K .....	64
Replace system management module - ASAA1K .....	68
ASAA70 and ASAA90 systems .....	73
Overview of the maintenance procedures - ASAA70 and ASAA90 .....	73
Boot media .....	74
Chassis .....	86
Controller .....	93
Replace a DIMM - ASAA70 and ASAA90 .....	109
Replace an SSD drive - ASAA70 and ASAA90 .....	115
Replace a fan module - ASAA70 and ASAA90 .....	118
Replace NVRAM - ASAA70 and ASAA90 .....	123
Replace the NV battery - ASAA70 and ASAA90 .....	130
I/O module .....	136
Hot-swap a power supply - ASAA70 and ASAA90 .....	151
Replace the real-time clock battery - ASAA70 and ASAA90 .....	154
Replace the system management module - ASAA70 and ASAA90 .....	160
ASAA20, ASAA30, and ASAA50 systems .....	166
Overview of hardware maintenance - ASAA20, ASAA30, and ASAA50 .....	166
Boot media .....	167
Chassis .....	181
Controller .....	189
Replace a DIMM - ASAA20, ASAA30, and ASAA50 .....	209
Replace a Drive - ASAA20, ASAA30, and ASAA50 .....	216
Replace a fan module - ASAA20, ASAA30, and ASAA50 .....	218
I/O module .....	224
Replace the NV battery - ASAA20, ASAA30, and ASAA50 .....	240
Hot-swap a power supply - ASAA20, ASAA30, and ASAA50 .....	246
Replace the real-time clock battery - ASAA20, ASAA30, and ASAA50 .....	250
ASA C30 systems .....	257
Overview of hardware maintenance - ASA C30 .....	257
Boot media .....	258

Chassis .....	272
Controller .....	280
Replace a DIMM - ASA C30 .....	300
Replace a drive - ASA C30 .....	307
Replace a fan module - ASA C30 .....	309
I/O module .....	315
Replace the NV battery - ASA C30 .....	331
Replace a power supply - ASA C30 .....	337
Replace the real-time clock battery - ASA C30 .....	341

# Maintain ASA r2 systems

## ASA A1K systems

### Overview of the maintenance procedures - ASA A1K

Maintain the hardware of your ASA A1K 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 ASA A1K system has already been deployed as a storage node in the ONTAP environment.

#### System components

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

<a href="#">Boot media - automated recovery</a>	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media.
<a href="#">Controller</a>	A controller consists of a board, firmware, and software. It controls the storage, I/O cards, 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</a>	The NVRAM (Non-Volatile Random Access Memory) is a module that allows the controller to protect and save in-flight data if the system loses power. The system ID resides in the NVRAM module. When replaced, the controller assumes the new system ID from the replacement NVRAM module.
<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.

## System management module

The System management module provides the interface between the controller and a console or laptop for controller or system maintenance purposes. The System management module contains the boot media and stores the system serial number (SSN).

## Boot media

### Boot media replacement workflow - ASA A1K

Get started with replacing the boot media in your ASA A1K storage system by reviewing the replacement requirements, shutting down the controller, replacing the boot media, restoring the image on the boot media, 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 - ASA A1K

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

Review the following requirements.

- You must replace the failed boot media with a replacement boot media you received from NetApp.
- The cluster ports are used to communicate between the two controllers during the automated boot recovery process. Make sure that the cluster ports on the impaired controller are working properly.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcards/kmip/servers.cfg
  - /cfcards/kmip/certs/client.crt
  - /cfcards/kmip/certs/client.key
  - /cfcards/kmip/certs/CA.pem
- Understand the controller terminology used in this procedure:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

## What's next

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

### Shut down the controller to replace the boot media - ASA A1K

Shut down the impaired controller in your ASA A1K 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, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

#### About this task

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

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

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

#### Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

## What's next

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

### Replace the boot media - ASA A1K

The boot media in your ASA A1K 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.

#### Steps

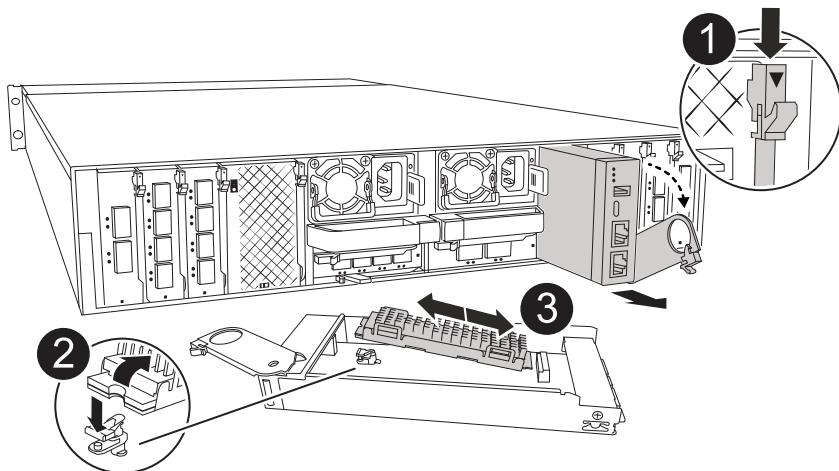
1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.



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

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

4. Remove the boot media from the management module:



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

- Press the blue locking button.
- 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:

- Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
- Rotate the boot media down toward the locking button.
- Push the locking button, rotate the boot media all the way down and then release the locking button.

6. Reinstall the System Management module:

- Align the module with the edges of the enclosure slot opening.
- 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.

- Reable the System Management module.

8. Plug the power cables into the power supplies and reinstall the power cable retainer.

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

### What's next

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

## Restore the ONTAP image on the boot media - ASA A1K

After installing the new boot media device in your ASA A1K system, you can start the automated boot media recovery process to restore the configuration from the partner node.

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

### Before you begin

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

### Steps

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

```
boot_recovery -partner
```

The screen displays the following message:

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

2. Monitor the boot media install recovery process.

The process completes and displays the Installation complete message.

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

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



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

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

## Onboard Key Manager (OKM)

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

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

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

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

### Show example of passphrase and backup data prompts

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

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

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

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

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

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

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

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

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

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

### External Key Manager (EKM)

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

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

a. Enter the EKM configuration settings when prompted:

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

#### Show example of client certificate contents

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

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

#### Show example of client key file contents

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

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

#### Show example of KMIP server file contents

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

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

**Show example of server configuration file contents**

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

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

**Show example of ONTAP Cluster UUID prompt**

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

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

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

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

### Show example of temporary network setting prompts

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

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

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

#### b. Verify the key restoration status:

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

### Show example of key recovery error and warning messages

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

Uptime: 11m32s
System halting...

LOADER-B>
```

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

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

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

5. If automatic giveback was disabled, reenable it:

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

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

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

#### What's next

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

#### Return the failed part to NetApp - ASA A1K

If a component in your ASA A1K system fails, return the failed part to NetApp. See the [Part Return and Replacements](#) page for further information.

# Chassis

## Chassis replacement workflow - ASA A1K

Get started with replacing the chassis of your ASA A1K storage system by reviewing the replacement requirements, shutting down the controller, replacing the chassis, and verifying system operations.

1

### Review the chassis replace requirements

Review the chassis replacement requirements.

2

### Prepare for chassis replace

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

3

### Shut down the controller

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

4

### Replace the chassis

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

5

### Complete the chassis replacement

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

## Requirements to replace the chassis - ASA A1K

Before replacing the chassis in your ASA A1K system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have local administrator credentials for ONTAP, the correct replacement chassis, and the necessary tools.

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

Review the following requirements.

- Make sure all other components in the system are functioning properly; if not, contact [NetApp support](#) for assistance.
- Obtain local administrator credentials for ONTAP if you don't have them.
- Make sure that you have the necessary tools and equipment for the replacement.
- You can use the chassis replacement procedure with all versions of ONTAP supported by your system.

- The chassis replacement procedure is written with the assumption that you are moving the bezel, fans, controller module, NVRAM12, system management module, I/O cards and blanking modules, and PSUs to the new chassis, and that the replacement chassis is a new component from NetApp.

## What's next?

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

## Prepare to replace the chassis - ASA A1K

Prepare to replace the impaired chassis in your ASA A1K 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 ASA A1K chassis hardware, you need to [shut down the controller](#).

### Shut down the controller to replace the chassis - ASA A1K

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

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

#### About this task

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

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

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

#### Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

## What's next?

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

### Replace the chassis - ASA A1K

Replace the chassis of your ASA A1K 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

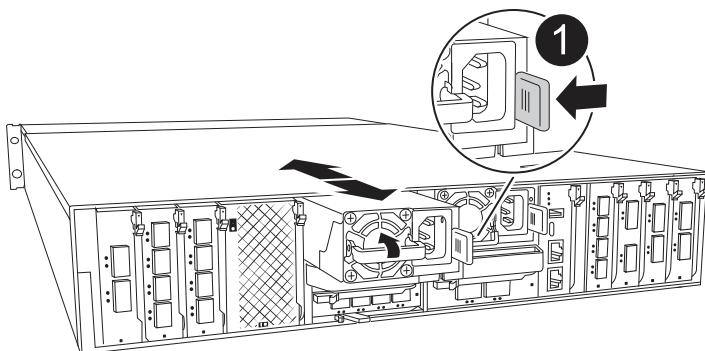
1. Remove the PSUs:
  - a. If you are not already grounded, properly ground yourself.
  - b. Unplug power cords from the PSUs.

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.

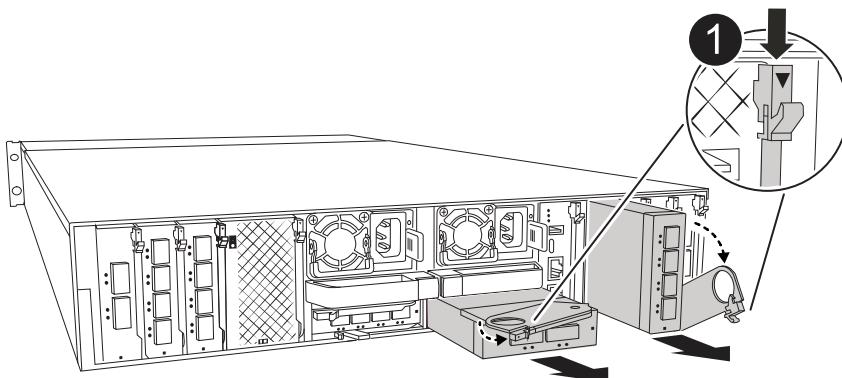


Cables should have been labeled at the beginning of this procedure.

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.

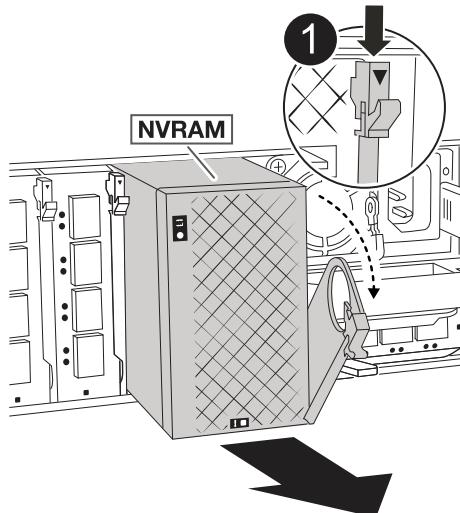
2. Remove the NVRAM12 module:

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.

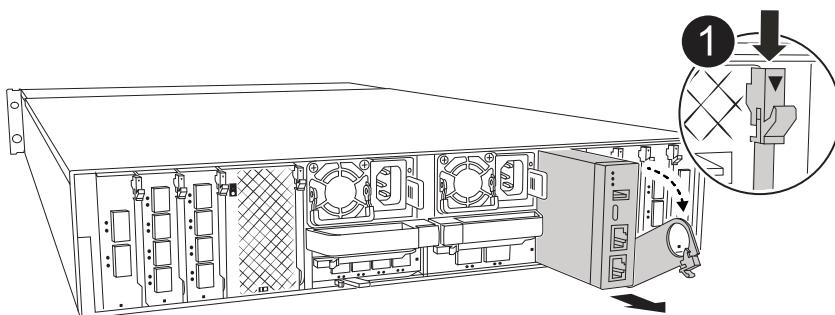


1	NVRAM12 cam latch
---	-------------------

d. Set the NVRAM module on a stable surface.

3. Remove the system management module:

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

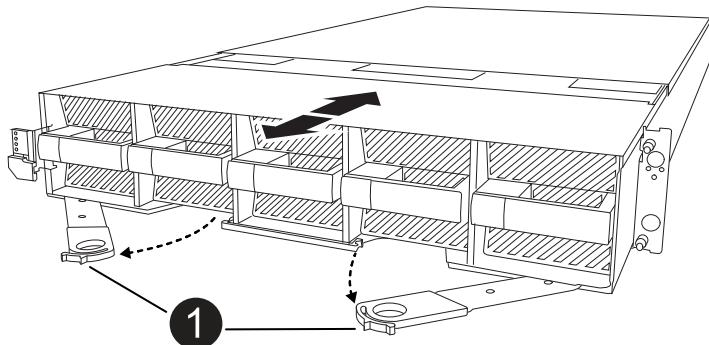


1	System Management module cam latch
---	------------------------------------

### Step 3: Remove the controller module

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

The controller module moves slightly out of the chassis.



1

Locking cam latches

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

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

#### Step 4: Replace the impaired chassis

Remove the impaired chassis and install the replacement chassis.

##### Steps

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

#### Step 5: Install the chassis components

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

##### Steps

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

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



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

Make sure that the cables are connected according to the cable labels.

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

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



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

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

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

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

## What's next?

After you've replaced the impaired ASA A1K chassis and reinstalled the components into it, you need to [complete the chassis replacement](#).

### Complete the chassis replacement - ASA A1K

Reboot the controller, verify system health, and return the failed part to NetApp to complete the final step in the ASA A1K chassis replacement procedure.

#### Step 1: Boot the controllers and verify system health

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

## Steps

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

#### Step 2: Return the failed part to NetApp

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

## Controller

### Controller replacement workflow - ASA A1K

Get started with replacing the controller in your ASA A1K storage system by shutting down the impaired controller, removing and replacing the controller, reconfiguring the system settings, and verifying system operations.

1

#### [Review controller replacement requirements](#)

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

2

#### [Shut down the impaired controller](#)

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

3

#### [Replace the controller](#)

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

4

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

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

5

### Give back the controller

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

6

### Complete controller replacement

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

## Requirements to replace the controller - ASA A1K

Before replacing the controller in your ASA A1K 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.

- 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 a replacement component you received from NetApp.
- Do not use this procedure for controller upgrades; instead, refer to the [Choose your controller hardware upgrade procedure](#) for guidance.
- You must replace a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- Because the boot device is located on the System Management module that is installed in the back of the system, you do not need to move the boot device when replacing a controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text log file.

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

## What's next?

After you've reviewed the requirements to replace your ASA A1K controller, you need to [shut down the controllers](#).

## Shut down the impaired controller - ASA A1K

Shut down the controller in your ASAA1K storage system to prevent data loss and ensure system stability when replacing the controller.

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

### About this task

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

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

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

### Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

## What's next?

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

### Replace the controller - ASA A1K

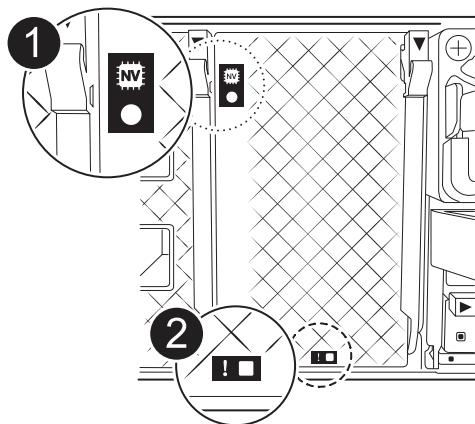
Replace the controller of your ASA A1K system when a hardware failure requires it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

#### Step 1: Remove the controller module

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

#### Steps

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

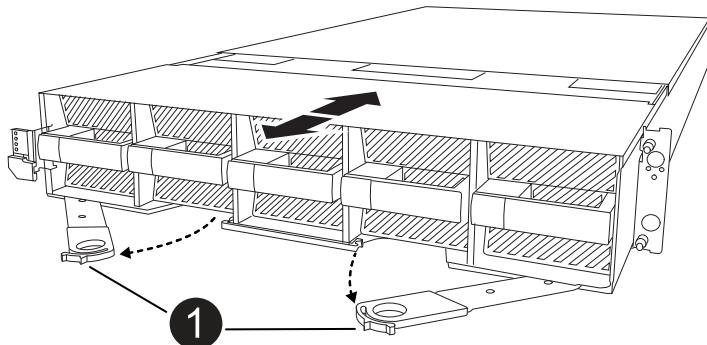


1	NVRAM status LED
2	NVRAM attention LED

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

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

The controller module moves slightly out of the enclosure.



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

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

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

#### Step 2: Move the fans

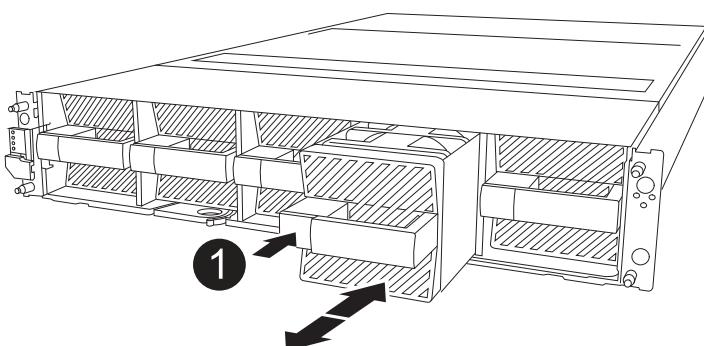
You must remove the five fan modules from the impaired controller module to the replacement controller module.

##### Steps

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



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



1	Black locking button
---	----------------------

4. Install the fan in the replacement controller module:

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

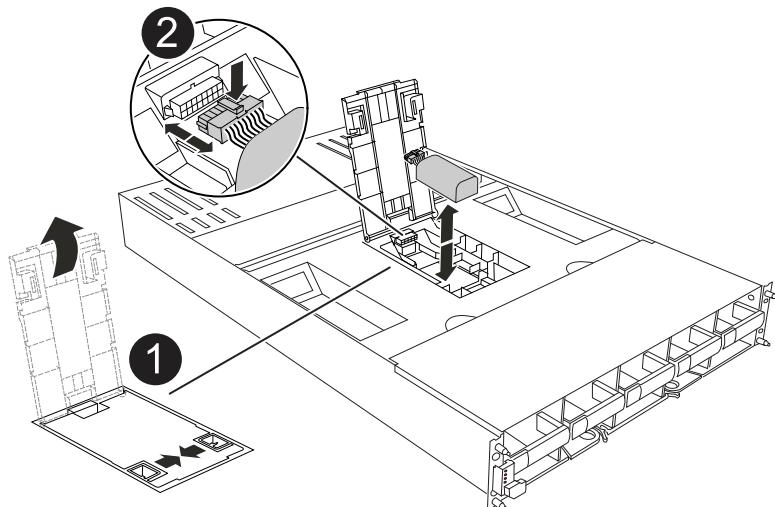
5. Repeat the preceding steps for the remaining fan modules.

#### Step 3: Move the NV battery

Move the NV battery to the replacement controller.

##### Steps

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



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

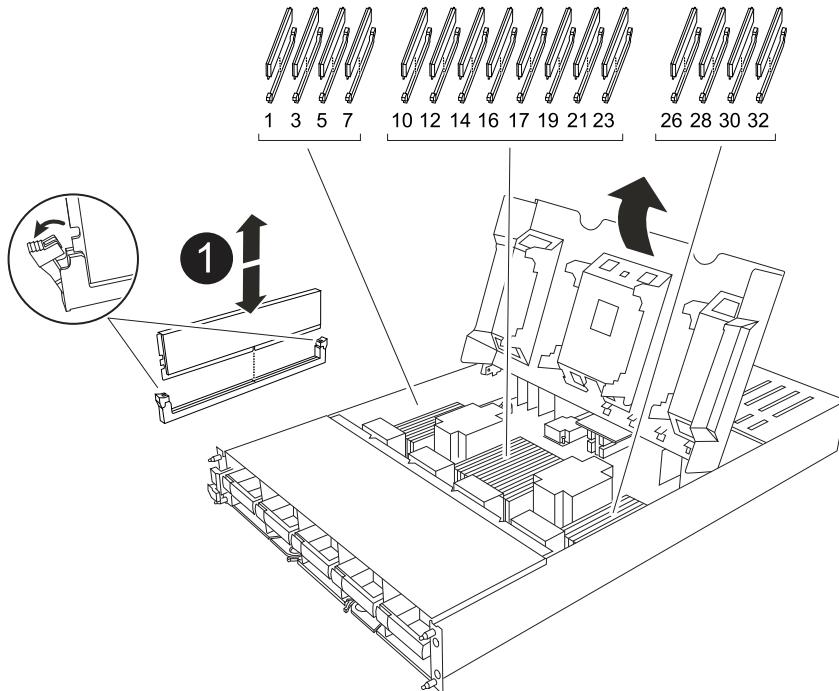
2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module.
5. Move the battery pack to the replacement controller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

#### Step 4: Move system DIMMs

Move the DIMMs to the replacement controller module.

##### Steps

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



1

System DIMM

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



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

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

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



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

6. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
7. Repeat these steps for the remaining DIMMs.

Close the motherboard air duct.

#### Step 5: Install the controller module

Reinstall the controller module and boot it.

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

It must lie flush against the controller module sheet metal.

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



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



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

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



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

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

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

## What's next?

After you've replaced the impaired ASA A1K controller, you need to [restore the system configuration](#).

### Restore and verify the system configuration - ASA A1K

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

#### Step 1: Verify HA config settings

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

#### Steps

1. Boot to maintenance mode: `boot_ontap maint`

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

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

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



If you see *PERSONALITY MISMATCH* contact customer support.

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

The HA state should be the same for all components.

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

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

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

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

## Step 2: Verify disk list

### Steps

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

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

2. Exit Maintenance mode: `halt`.

### What's next?

After you've restored and verified the system configuration for your ASA A1K system, you need to [give back the controller](#).

### Give back the controller - ASA A1K

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

## No encryption

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

### Steps

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

## Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

### Steps

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



You are prompted twice for the passphrase.

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



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

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

The output should show no results when when properly synchronized.

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

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

## What's next?

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

## Complete controller replacement - ASA A1K

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

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

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

#### Steps

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

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

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

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

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

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

### Step 2: Return the failed part to NetApp

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

## Replace a DIMM - ASA A1K

Replace a DIMM in your ASA A1K system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

### Before you begin

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

### Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

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

### About this task

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

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

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

### Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

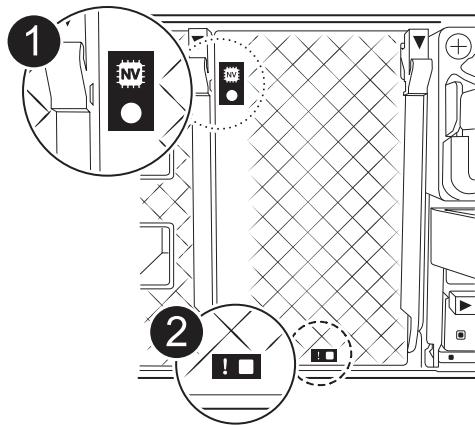
If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover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

## Step 2: Remove the controller module

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

### Steps

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

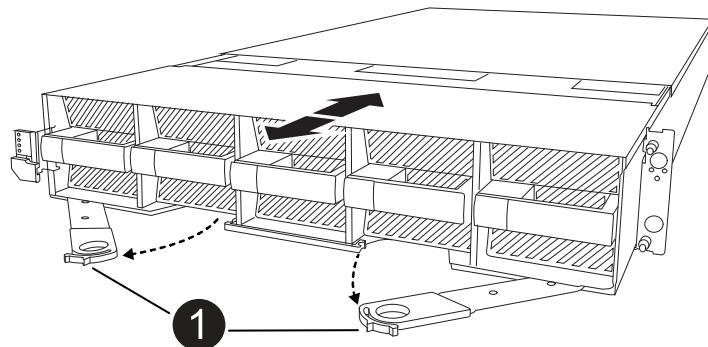


1	NVRAM status LED
2	NVRAM attention LED

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

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

The controller module moves slightly out of the enclosure.



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

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

### Step 3: Replace a DIMM

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

#### Steps

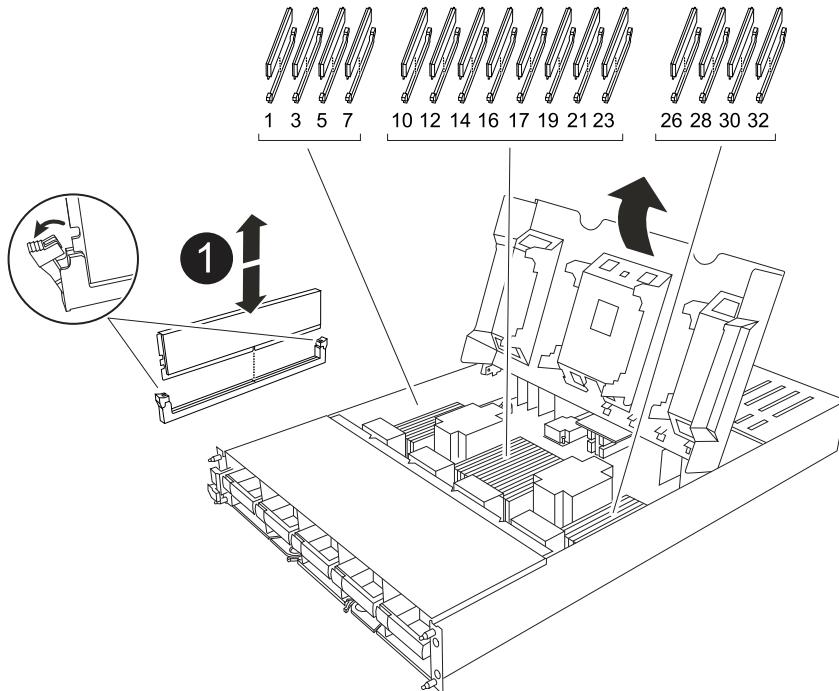
1. If you are not already grounded, properly ground yourself.
2. Open the controller air duct on the top of the controller.
  - a. Insert your fingers in the recesses at the far ends of the air duct.
  - b. Lift the air duct and rotate it upward as far as it will go.
3. Locate the DIMMs on your controller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

##### Steps

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

It must lie flush against the controller module sheet metal.

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

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



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

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

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

## Step 5: Return the failed part to NetApp

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

## Replace a fan - ASA A1K

Replace a failed or faulty fan module in your ASA A1K 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.

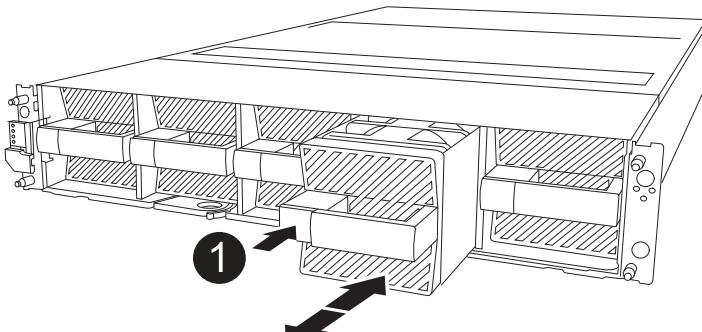


There is a single LED for each fan. It is green when the fan is functioning correctly and amber when not.

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.



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

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

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

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

## Replace NVRAM - ASA A1K

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

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

### Before you begin

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

### Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

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

### About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv

advanced mode) displays the node name, [quorum status](#) of that node, availability status of that node, and operational status of that node.

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

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

## Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

## Step 2: Replace the NVRAM module or NVRAM DIMM

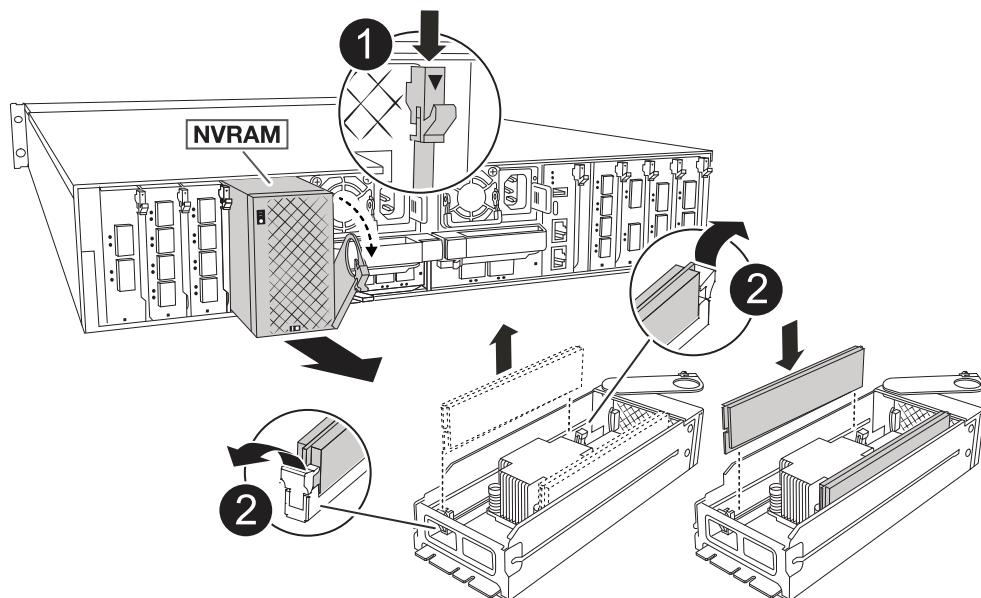
Replace the NVRAM module or NVRAM DIMMs using the appropriate following option.

## Option 1: Replace the NVRAM module

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

### Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.
3. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the tray down.
4. Remove the impaired NVRAM module from the enclosure:
  - a. Depress the locking cam button.
  - b. Rotate the cam latch down as far as it will go.
  - c. Remove the impaired NVRAM module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.



1	Cam locking button
2	DIMM locking tabs

5. Set the NVRAM module on a stable surface.
6. Remove the DIMMs, one at a time, from the impaired NVRAM module and install them in the replacement NVRAM module.
7. Install the replacement NVRAM module into the enclosure:
  - a. Align the module with the edges of the enclosure opening in slot 4/5.
  - b. Gently slide the module into the slot all the way, and then rotate the cam latch all the way up to lock the module in place.

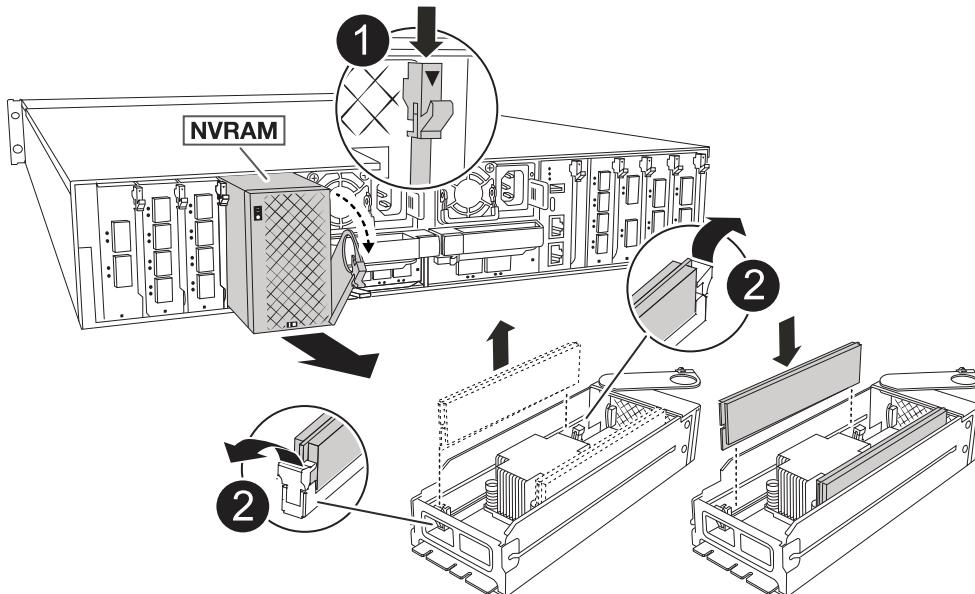
8. Recable the PSUs.
9. Rotate the cable management tray up to the closed position.

#### Option 2: Replace the NVRAM DIMM

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

#### Steps

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



<b>1</b>	Cam locking button
<b>2</b>	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 component, you must reboot the controller module by plugging the power cables back into the PSU.

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

### Step 4: Verify controller state

You must confirm the controller state of the controllers connected to the disk pool when you boot the controller.

#### Steps

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

In the command output, you should see a message indicates the state of the controllers.

Node	Partner	Takeover	Possible State Description
-----	-----	-----	-----
-----	-----	-----	-----
<nodename>			
<nodename>- P2-3-178			
true			
Connected to <nodename>-P2-3-178.			
P2-3-178			
Waiting for cluster applications			
to			
come online on the local node.			
AFF-A90-NBC-P2-3-178			
<nodename>- P2-3-177			
true			
Connected to <nodename>-P2-3-177,			
P2-3-177			
Partial giveback			
2 entries were displayed.			

4. Give back the controller:

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

The controller connects back its storage pool and completes booting.

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



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

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

- b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`
5. Verify all disks are displayed: `storage disk show`

```
::> storage disk show
      Usable          Disk      Container  Container
Disk      Size Shelf Bay Type      Type      Name
-----
1.0.0      3.49TB    0  SSD-NVM aggregate  pod_NVME_SSD_1
1.0.1      3.49TB    1  SSD-NVM aggregate  pod_NVME_SSD_1
1.0.2      3.49TB    2  SSD-NVM aggregate  pod_NVME_SSD_1
1.0.3      3.49TB    3  SSD-NVM aggregate  pod_NVME_SSD_1
1.0.4      3.49TB    4  SSD-NVM aggregate  pod_NVME_SSD_1
[...]
48 entries were displayed.
```

## 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 - ASA A1K

Replace the NV battery in your ASA A1K 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, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

### About this task

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

operational status of that node.

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

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

## Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

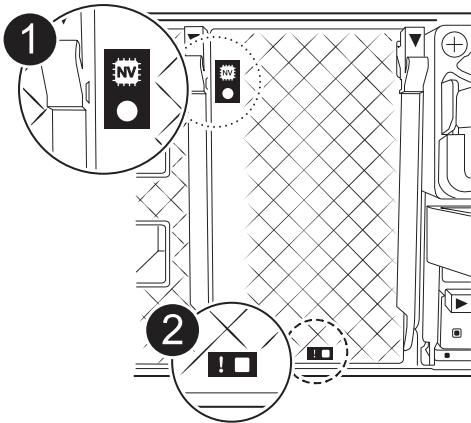
If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller:  <pre>storage failover takeover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

## Step 2: Remove the controller module

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

## Steps

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

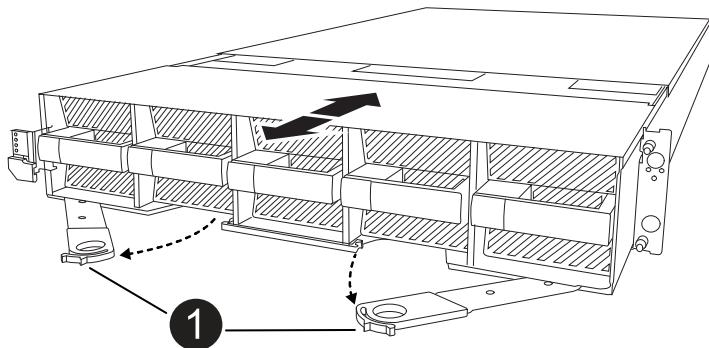


1	NVRAM status LED
2	NVRAM attention LED

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

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

The controller module moves slightly out of the enclosure.



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

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

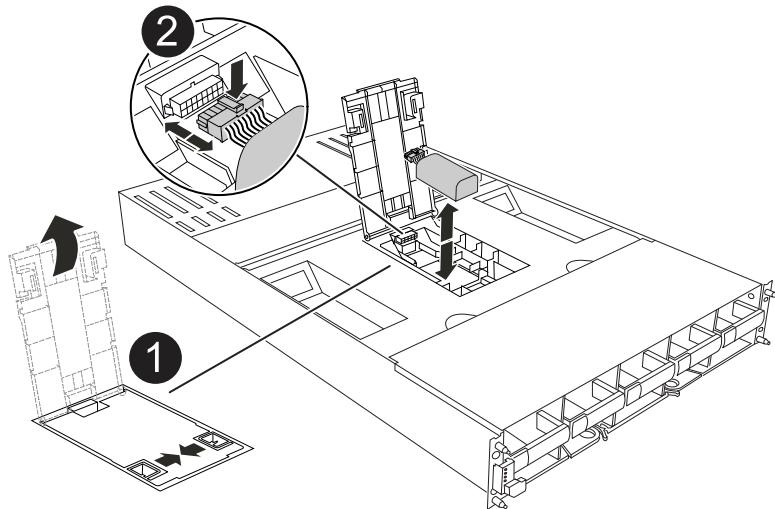
Make sure that you support the bottom of the controller module as you slide it out of the enclosure.

### Step 3: Replace the NV battery

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

## Steps

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



1	NV battery air duct cover
2	NV battery plug

2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module, and then set it aside.
5. Remove the replacement battery from its package.
6. Install the replacement battery pack into the controller:
  - a. Plug the battery plug into the riser socket and make sure that the plug locks into place.
  - b. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
7. Close the NV air duct cover.

Make sure that the plug locks into the socket.

## Step 4: Reinstall the controller module

Reinstall the controller module and boot it.

## Steps

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

It must lie flush against the controller module sheet metal.
2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.

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



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

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

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

## Step 5: Return the failed part to NetApp

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

## I/O module

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

The ASA A1K 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 ASA A1K storage system with the same type of I/O module, or with a different kind of I/O module. You can also add an I/O module into a system with empty slots.

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

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

- [Hot swap an I/O module](#)

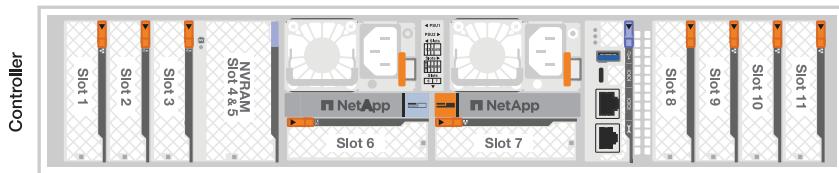
Hot swapping an I/O module allows you to replace a failed module without shutting down the system, minimizing downtime and maintaining system availability.

- [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 ASA A1K controllers are numbered 1 through 11, as shown in the following illustration.



## Add an I/O module - ASA A1K

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

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

### About this task

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

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

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

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

Shut down or take over 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.

### About this task

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

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

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

### Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

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

#### Before you begin

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

## Add I/O module to an available slot

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

### Steps

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



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

6. Rotate the cable management tray up to the closed position.

7. From the LOADER prompt, reboot the node:

bye



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

8. Give back the controller from the partner controller:

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

9. Repeat these steps for controller B.

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

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

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

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

## Add I/O module to a fully-populated system

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

### About this task

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

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See <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.



If you encounter an issue during reboot, see [BURT 1494308 - Environment shutdown might be triggered during I/O module replacement](#)

10. Give back the controller from the partner controller:

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

11. Enable automatic giveback if it was disabled:

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

12. Do one of the following:

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

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

- If you removed a NIC I/O module and installed a storage I/O module, install and cable your NS224 shelves, as described in [Hot-add workflow](#).

13. Repeat these steps for controller B.

## Hot swap an I/O module - ASA A1K

You can hot swap an Ethernet I/O module in your ASA A1K 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 <i>y</i> 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: Replace the failed I/O module**

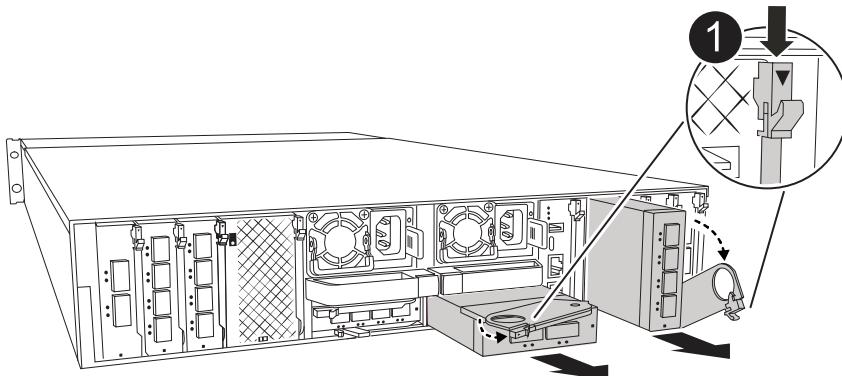
Replace the failed I/O module with an equivalent I/O module.

**Steps**

1. If you are not already grounded, properly ground yourself.
2. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
3. Remove the I/O module from the controller module:



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



1

Cam locking button

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

Keep track of which slot the I/O module was in.

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

#### Step 4: Bring the replacement I/O module online

Bring the replacement I/O module online, verify the I/O module ports initialized successfully, verify the slot is powered on, and then verify the I/O module is online and recognized.

#### About this task

After the I/O module is replaced and the ports are returned to a healthy state, LIFs are reverted to the replaced I/O module.

#### Steps

1. Bring the replacement I/O module online:

a. Enter the following command:

```
system controller slot module insert -node impaired_node_name -slot slot_number
```

b. Enter *y* when you see the prompt, *Do you want to continue?*

The output should confirm the I/O module was successfully brought online (powered on, initialized, and placed into service).

For example, the following command brings slot 7 on node 2 (the impaired controller) online, and displays a message that the process was successful:

```
node2::> system controller slot module insert -node node2 -slot 7

Warning: IO_2X_100GBE_NVDA_NIC module in slot 7 of node node2 will be
powered on and initialized.

Do you want to continue? {y|n}: `y'

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

2. Verify that each port on the I/O module successfully initialized:

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

```
event log show -event *hotplug.init*
```



It might take several minutes for any required firmware updates and port initialization.

The output should show one or more hotplug.init.success EMS events and *hotplug.init.success*: in the *Event* column, indicating each port on the I/O module initialized successfully.

For example, the following output shows initialization succeeded for I/O ports e7b and e7a:

```
node2::> event log show -event *hotplug.init*
Time           Node           Severity      Event
-----
-----
7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e7b" in slot 7 succeeded

7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e7a" in slot 7 succeeded

2 entries were displayed.
```

b. If the port initialization fails, review the EMS log for the next steps to take.

3. Verify the I/O module slot is powered on and ready for operation:

```
system controller slot module show
```

The output should show the slot status as *powered-on* and therefore ready for operation of the I/O module.

4. Verify that the I/O module is online and recognized.

Enter the command from the console of the impaired controller:

```
system controller config show -node local -slot slot_number
```

If the I/O module was successfully brought online and is recognized, the output shows I/O module information, including port information for the slot.

For example, you should see output similar to the following for an I/O module in slot 7:

```

node2::> system controller config show -node local -slot 7

Node: node2
Sub- Device/
Slot slot Information
-----
7 - Dual 40G/100G Ethernet Controller CX6-DX
    e7a MAC Address: d0:39:ea:59:69:74 (auto-100g_cr4-fd-
up)
        QSFP Vendor: CISCO-BIZLINK
        QSFP Part Number: L45593-D218-D10
        QSFP Serial Number: LCC2807GJFM-B
    e7b MAC Address: d0:39:ea:59:69:75 (auto-100g_cr4-fd-
up)
        QSFP Vendor: CISCO-BIZLINK
        QSFP Part Number: L45593-D218-D10
        QSFP Serial Number: LCC2809G26F-A
        Device Type: CX6-DX PSID(NAP0000000027)
        Firmware Version: 22.44.1700
        Part Number: 111-05341
        Hardware Revision: 20
        Serial Number: 032403001370

```

#### Step 5: Restore the storage system to normal operation

Restore your storage system to normal operation by giving back storage to the controller that was taken over (as needed), restoring automatic giveback (as needed), verifying LIFs are on their home ports, and reenabling AutoSupport automatic case creation.

#### Steps

1. As needed for the version of ONTAP your storage system is running and the state of the controllers, give back storage and restore automatic giveback on the controller that was taken over:

If...	Then...
If either controller took over its partner automatically	<ol style="list-style-type: none"> <li>1. Return the controller that was taken over to normal operation by giving back its storage:           <pre>storage failover giveback -ofnode controller that was taken over_name</pre> </li> <li>2. Restore automatic giveback from the console of the controller that was taken over:           <pre>storage failover modify -node local -auto -giveback true</pre> </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 an I/O module - ASA A1K

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

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

#### Before you begin

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

#### Step 1: Shut down the impaired node

Shut down or take over the impaired controller.

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

#### About this task

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

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

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a

healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

## Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

## Step 2: Replace a failed I/O module

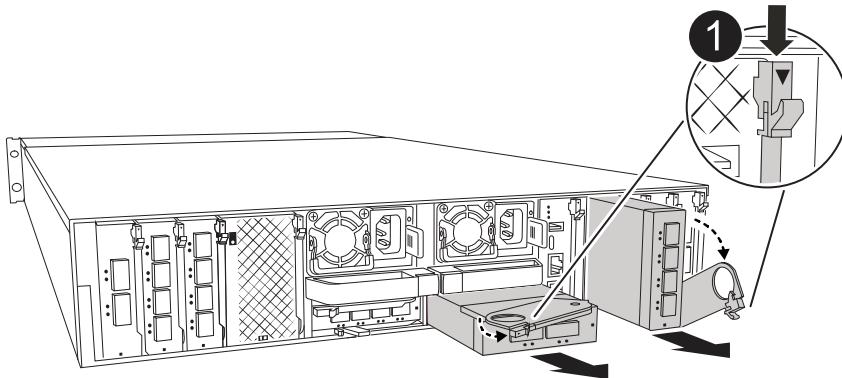
To replace an I/O module, locate it within the enclosure and follow the specific sequence of steps.

## Steps

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



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



1

I/O cam latch

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

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

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

### Step 3: Reboot the controller

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

#### Steps

1. Reboot the controller from the LOADER prompt:

bye



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

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

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

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

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

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

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

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

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

## Hot-swap a power supply - ASA A1K

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

The power supplies are redundant and hot-swappable. You do not have to shut down the controller to replace a PSU.

### About this task

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



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

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

## Option 1: Hot-swap an AC PSU

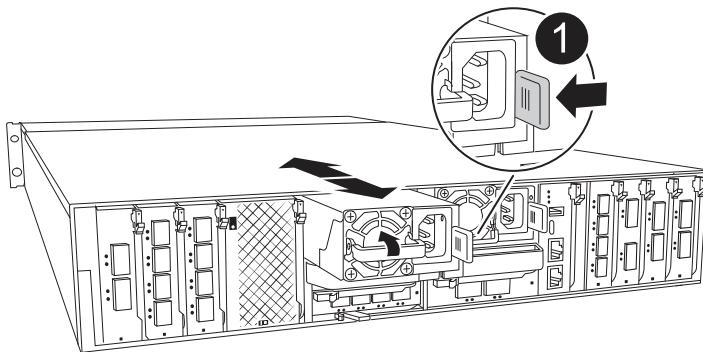
To replace an AC PSU, complete the following steps.

### Steps

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



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



Terracotta PSU locking tab

5. Install the replacement PSU in the controller module:

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

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



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

6. Reconnect the PSU cabling:

- a. Reconnect the power cable to the PSU.
- b. Secure the power cable to the PSU using the power cable retainer.

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

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

### Option 2: Hot-swap a DC PSU

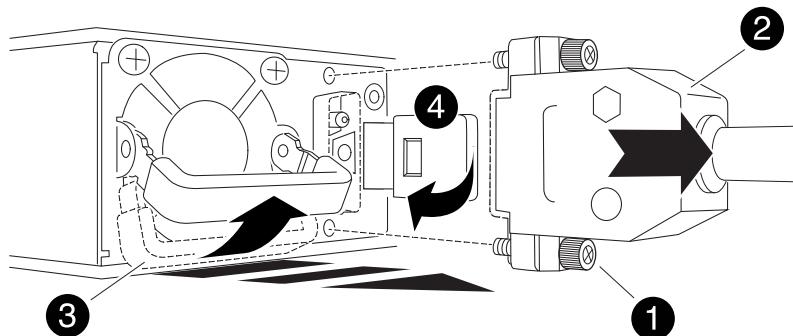
To replace a DC PSU, complete the following steps.

#### Steps

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



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



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

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

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

way.



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

6. Reconnect the D-SUB DC power cable:

- a. Plug the power cable connector into the PSU.
- b. Secure the power cable to the PSU with the thumbscrews.

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

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

## Replace the real-time clock battery - ASA A1K

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your ASA A1K 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 [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, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

### About this task

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

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

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

### Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

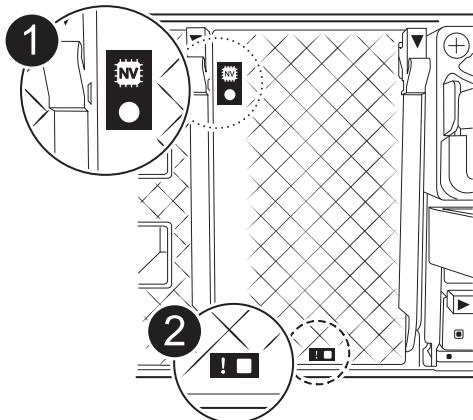
If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: <pre>storage failover takeover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

## Step 2: Remove the controller module

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

### Steps

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

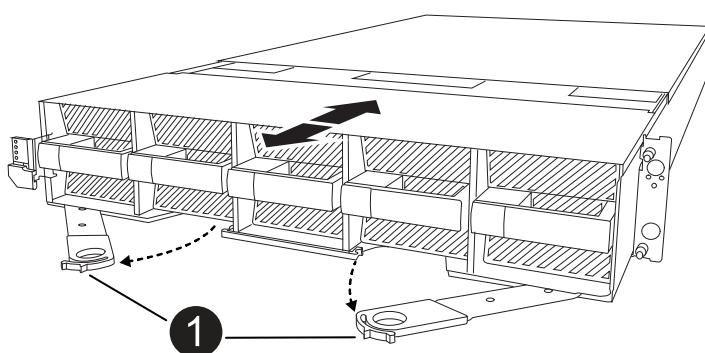


1	NVRAM status LED
2	NVRAM attention LED

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

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

The controller module moves slightly out of the enclosure.



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

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

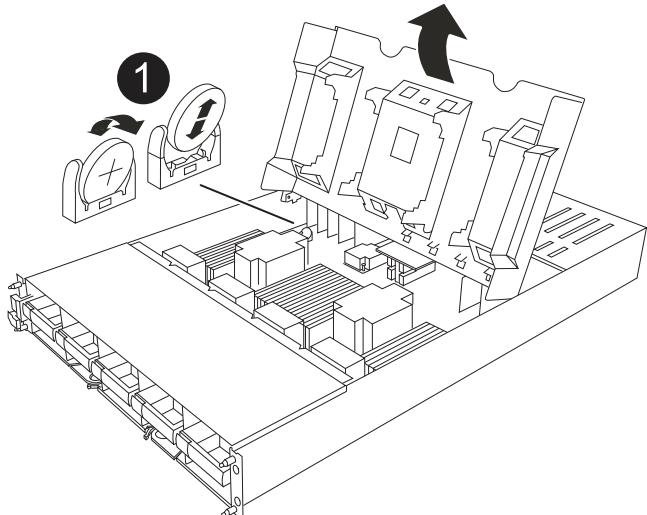
Make sure that you support the bottom of the controller module as you slide it out of the enclosure.

### Step 3: Replace the RTC battery

Remove failed RTC battery and install the replacement RTC battery.

#### Steps

1. Open the controller air duct on the top of the controller.
  - a. Insert your fingers in the recesses at the far ends of the air duct.
  - b. Lift the air duct and rotate it upward as far as it will go.
2. Locate the RTC battery under the air duct.



1

RTC battery and housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.



Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

#### Step 4: Reinstall the controller module

Reinstall the controller module and boot it.

##### Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.  
It must lie flush against the controller module sheet metal.
2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans



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

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

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

## Step 5: Reset the time and date on the controller

After replacing the RTC battery, inserting controller and powering on first BIOS reset, you will see the following error messages:



RTC date/time error. Reset date/time to default

RTC power failure error

These messages are expected and you can continue with this procedure.

### Steps

1. Check the date and time on the healthy controller with the `cluster date show` command.



If your system stops at the boot menu, select the option for `Reboot node` and respond `y` when prompted, then boot to LOADER by pressing `Ctrl-C`

- a. At the LOADER prompt on the target controller, check the time and date with the `cluster date show` command.
- b. If necessary, modify the date with the `set date mm/dd/yyyy` command.
- c. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
  1. Confirm the date and time on the target controller.
  2. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.

## Step 6: Return the failed part to NetApp

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

## Replace system management module - ASA A1K

Replace the System Management module in your ASA A1K system when it becomes defective or its firmware is corrupted. The replacement process involves shutting down the controller, replacing the failed System Management module, rebooting the controller, updating the license keys, and returning the failed part to NetApp.

The System Management module, located at the back of the controller in slot 8, contains onboard components for system management, as well as ports for external management. The target controller must be shut down to replace an impaired System Management module or replace the boot media.

The System Management module has the following onboard components:

- Boot media, allowing boot media replacement without removing the controller module.

- BMC
- Management switch

The System Management module also contains the following ports for external management:

- RJ45 Serial
- USB Serial (Type-C)
- USB Type-A (Boot recovery)
- e0M RJ45 Ethernet

### Before you begin

- All other system components must be working properly.
- The partner controller must be able to take over the impaired controller.
- You must replace the failed component with a replacement FRU component you received from your provider.

### 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, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

### About this task

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

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

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

### Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

## Step 2: Replace the impaired System Management module

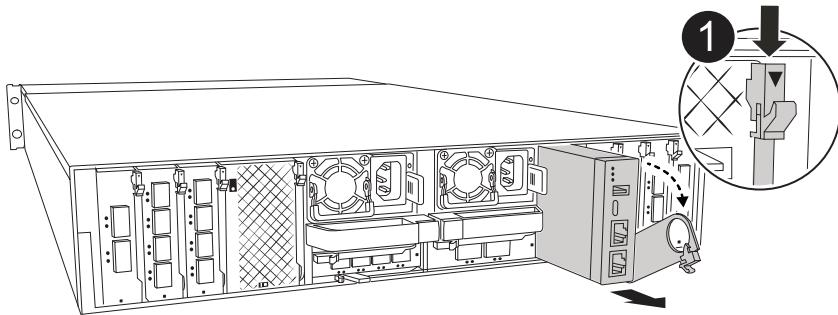
Replace the impaired system management module.

### Steps

- 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

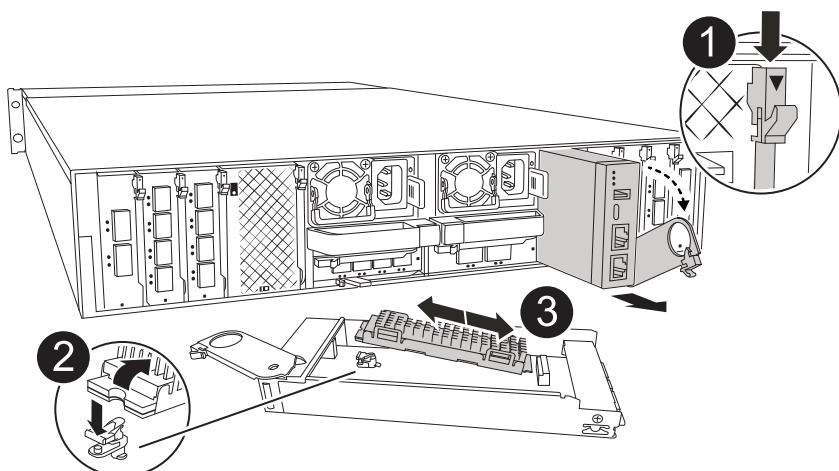
System Management module cam latch

- If you are not already grounded, properly ground yourself.
- Remove any cables connected to the System Management module. Make sure that label where the cables were connected, so that you can connect them to the correct ports when you reinstall the

module.

- c. Unplug the power supply cables from the PSUs.
- 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, so that the boot media is accessible.

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.

#### Steps

1. Plug the power cables back into the PSU.

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

2. Enter *bye* at the LOADER prompt.
3. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode _impaired_node_name_`
4. Restore automatic giveback by using the `storage failover modify -node local -auto -giveback true` command.
5. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

### Step 4: Install licenses and register serial number

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

#### About this task

Until you install license keys, features requiring standard licenses continue to be available to the node. However, if the node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed. Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the node as soon as possible.

#### Before you begin

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

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

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

#### Steps

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



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

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
  - a. Check for unused licenses: `license clean-up -unused -simulate`
  - b. If the list looks correct, remove the unused licenses: `license clean-up -unused`
4. Register the system serial number with NetApp Support.
  - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
  - If AutoSupport is not enabled, call [NetApp Support](#) to register the serial number.

#### Step 5: Return the failed part to NetApp

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

## ASA A70 and ASA A90 systems

### Overview of the maintenance procedures - ASA A70 and ASA A90

Maintain the hardware of your ASA A70 and ASA A90 storage systems 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 ASA A70 and ASA A90 systems have already been deployed as a storage node in the ONTAP environment.

#### System components

For the ASA A70 and ASA A90 storage systems, 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 storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media.
<a href="#">Chassis</a>	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
<a href="#">Controller</a>	A controller consists of a board, firmware, and software. It controls the storage, I/O cards, and runs the ONTAP operating system software.

DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
Drive	A drive is a device that provides the physical storage needed for data.
Fan	A fan cools the controller.
NVRAM	The NVRAM (Non-Volatile Random Access Memory) is a module that allows the controller to protect and save in-flight data if the system loses power. The system ID resides in the NVRAM module. When replaced, the controller assumes the new system ID from the replacement NVRAM module.
NV battery	The NV battery is responsible for providing power to the NVRAM module while data in-flight is being destaged to flash memory after a power loss.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.
System Management module	The System Management module provides the interface between the controller and a console or laptop for controller or system maintenance purposes. The System management module contains the boot media and stores the system serial number (SSN).

## Boot media

### Boot media replacement workflow - ASA A70 and ASA A90

Get started with replacing the boot media in your ASA A70 and ASA A90 storage system by reviewing the replacement requirements, shutting down the controller, replacing the boot media, restoring the image on the boot media, 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 - ASA A70 and ASA A90

Before replacing the boot media in your ASA A70 or ASA A90 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming the cluster ports on the impaired controller are working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

Review the following requirements.

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

## What's next

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

## Shut down the controller to replace the boot media - ASA A70 or ASA A90

Shut down the impaired controller in your ASA A70 or ASA A90 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, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

## About this task

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

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

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

## Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

## What's next

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

## Replace the boot media - ASA A70 and ASA A90

The boot media in your ASA A70 or ASA A90 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.

Replace the boot media, which is located inside the System Management module at the rear of the controller.

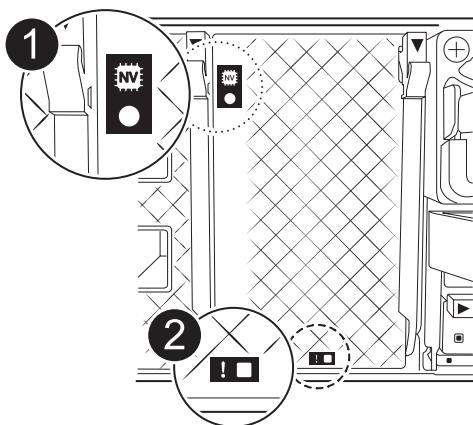
#### Before you begin

- You must have a replacement boot media.
- Have an anti-static mat available for the System Management module.

#### Steps

1. Verify that 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 NetApp Support for assistance.



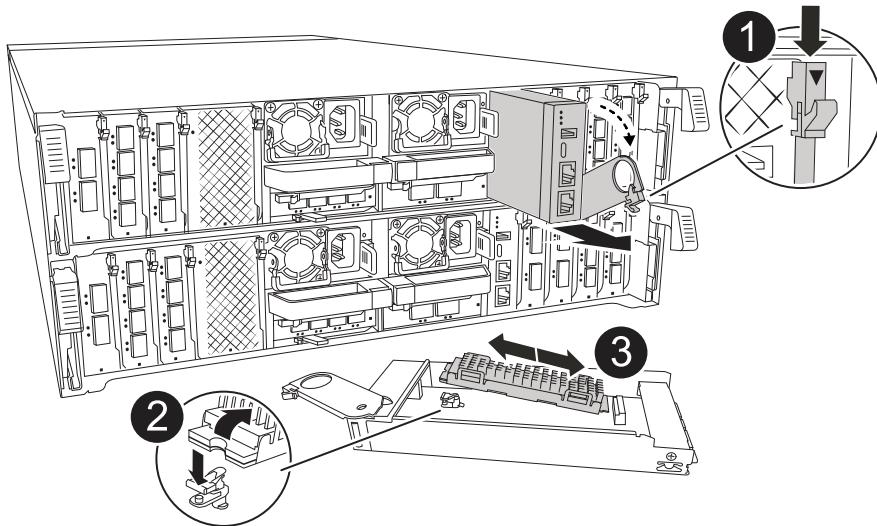
1	NVRAM status LED
2	NVRAM attention LED

2. Go to the rear of the chassis and properly ground yourself if you are not already grounded.
3. Disconnect power from the controller:
  - For AC power supplies, disconnect the power cords from the power supplies.
  - For DC power supplies, disconnect the power block from the power supplies.
4. Remove the System Management module:
  - a. Remove any cables connected to the System Management module. Label the cables to identify their correct ports for reinstallation.
  - b. Rotate the cable management arm down by pulling the buttons on both sides of the cable management arm.
  - c. Depress the system management cam button.

The cam lever moves away from the chassis.

- d. Rotate the cam lever all the way down and remove the System Management module from the controller.
- e. Place the System Management module on an anti-static mat with the boot media accessible.

5. Remove the boot media from the System Management module:



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

- a. Press the blue locking button.
- b. Rotate the boot media up, slide it out of the socket, and set it aside.

6. 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.

7. Reinstall the System Management module:

- a. Align the edges of the System Management module with the chassis opening.
- b. Gently slide the module into the chassis until the cam latch begins to engage.
- c. Rotate the cam latch all the way up to lock the module in place.
- d. Reconnect the cables to the System Management module using the labels you created during removal.
- e. Rotate the cable management arm up to the closed position.

8. Reconnect power to the controller:

- For AC power supplies, plug the power cords into the power supplies.
- For DC power supplies, reconnect the power block to the power supplies.

The controller automatically reboots when power is restored.

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

#### What's next

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

#### Restore the ONTAP image on the boot media - ASA A70 and ASA A90

After installing the new boot media device in your ASA A70 or ASA A90 system, you can start the automated boot media recovery process to restore the configuration from the partner node.

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

#### Before you begin

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

#### Steps

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

```
boot_recovery -partner
```

The screen displays the following message:

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

2. Monitor the boot media install recovery process.

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

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

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



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

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

## Onboard Key Manager (OKM)

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

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

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

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

### Show example of passphrase and backup data prompts

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

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

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

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

```
Successfully recovered keymanager secrets.
```

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

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

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

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

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

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

### External Key Manager (EKM)

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

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

a. Enter the EKM configuration settings when prompted:

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

#### Show example of client certificate contents

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

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

#### Show example of client key file contents

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

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

#### Show example of KMIP server file contents

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

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

**Show example of server configuration file contents**

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

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

**Show example of ONTAP Cluster UUID prompt**

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

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

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

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

### Show example of temporary network setting prompts

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

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

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

#### b. Verify the key restoration status:

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

### Show example of key recovery error and warning messages

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

Uptime: 11m32s
System halting...

LOADER-B>
```

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

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

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

5. If automatic giveback was disabled, reenable it:

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

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

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

#### What's next

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

#### **Return the failed part to NetApp - ASA A70 and ASA A90**

If a component in your ASA 70-90 system fails, return the failed part to NetApp. See the [Part Return and Replacements](#) page for further information.

# Chassis

## Chassis replacement workflow - ASA A70 and ASA A90

Get started with replacing the chassis in your ASA A70 or ASA A90 storage system by shutting down the controllers, replacing the chassis, and verifying system operations.

1

### Review the chassis replacement requirements

Review the chassis replacement requirements, including system compatibility, required tools, ONTAP credentials, and component functionality verification.

2

### Prepare for the chassis replacement

Prepare for the chassis replacement by locating the system, gathering credentials and tools, verifying the replacement chassis, and labeling cables.

3

### Shut down the controllers

Shut down the controllers to perform chassis maintenance safely.

4

### Replace the chassis

Move the components from the impaired chassis to the replacement chassis.

5

### Complete the chassis replacement

Complete the replacement by booting the controllers, performing giveback, and returning the failed chassis to NetApp.

## Requirement to replace the chassis - ASA A70 and ASA A90

Before replacing the chassis in your ASA A70 or ASA A90 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 enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

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, NVMe drives, and controller modules to the new chassis, and that the replacement chassis is a new component from NetApp.
- The chassis replacement procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

### What's next?

After reviewing the requirements, [prepare to replace the chassis](#).

### Prepare to replace chassis - ASA A70 and ASA A90

Prepare to replace the impaired chassis in your ASA A70 or ASA A90 system by identifying the impaired chassis, verifying the replacement components, and labeling the cables and controller modules.

#### Steps

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

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

 The Location LED remains lit for 30 minutes.
3. Before opening the packaging, examine the packaging label and verify the following:
  - Component part number
  - Part description
  - Quantity in the box
4. Remove the contents from the packaging and save the packaging for returning the failed component to NetApp.
5. Label all cables connected to the storage system. This ensures proper recabling later in this procedure.
6. Ground yourself if not already grounded.

### What's next?

After preparing to replace the ASA A70 or ASA A90 chassis, [shut down the controllers](#).

### Shut down the controllers to replace the chassis - ASA A70 and ASA A90

Shut down the controllers in your ASA A70 or ASA A90 storage system to prevent data loss and ensure system stability when replacing the chassis.

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

## Before you begin

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

## Steps

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

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

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

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

6. Exit the cluster shell:

```
exit
```

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

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

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

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



For clusters using SnapMirror synchronous operating in StrictSync mode: 

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

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

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

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

### What's next?

After the controllers are shut down, you need to [replace the chassis](#).

### Move and replace hardware - ASA A70 and ASA A90

Replace the chassis in your ASA A70 and ASA A90 when a hardware failure requires it. Remove the controllers and power supply units (PSUs), remove the drives, install the replacement chassis, and reinstall all components.

#### Step 1: Remove the PSUs and cables

You need to remove all four power supply units (PSUs), two per controller, before removing the controller. Removing them lightens the overall weight of each controller.

#### Steps

1. Remove the four PSUs:

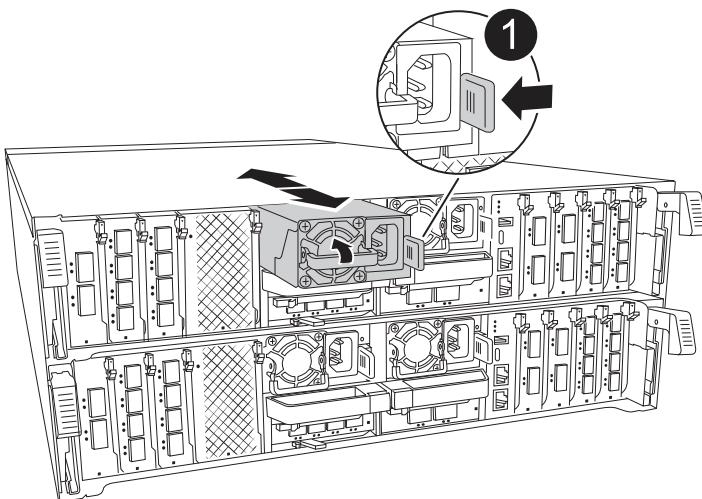
- If you are not already grounded, properly ground yourself.
- Unplug power cords from the controller module PSU.

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

- Remove the PSU from the controller 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 controller module.



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



Terracotta PSU locking tab

- Repeat these steps for the remaining PSUs.

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.

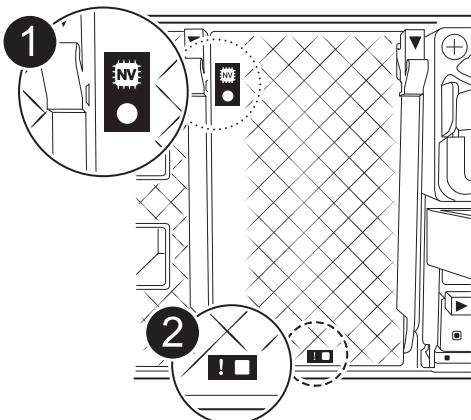
- b. Remove the cable management device from the controller modules and set them aside.

## Step 2: Remove the controller modules and drives

Remove the controllers from the chassis and then remove the drives from the chassis.

### Steps

1. Check the that amber NVRAM status LED located in slot 4/5 on the back of each controller module is off. Look for the NV icon.



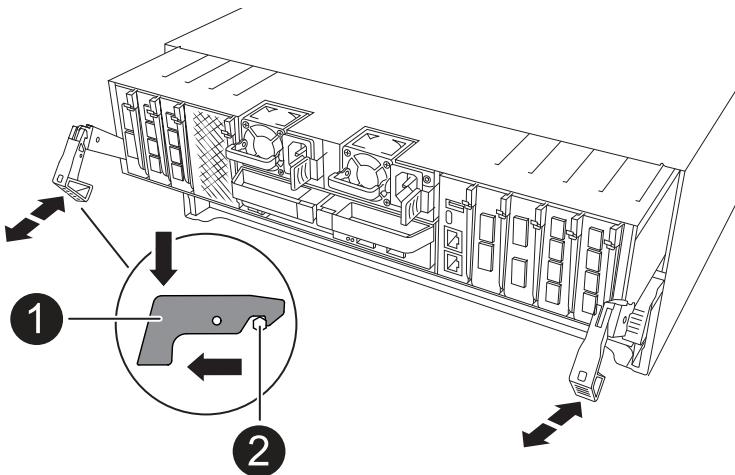
1	NVRAM status LED
2	NVRAM attention LED

- If the NVRAM LED is off, go to the next step.
- If the NVRAM LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact [NetApp Support Site](#) for assistance.

2. Remove the controller modules:

- a. Press down on both of the locking latches on the controller, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.



1	Locking latch
2	Locking pin

b. Slide the controller module out of the chassis by the locking latches, 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.

c. Repeat these steps for the second controller module.

3. Remove the drives:

- Gently remove the bezel from the front of the system.
- Press the release button at the top of the drive carrier face below the LEDs.
- Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

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



Drives are fragile. Always use two hands to support the drive weight when removing a drive to prevent damage to them.

- Keep track of which drive bay each drive was in and set the drive aside on a static-free cart or table.
- Repeat this step for the remaining drives in the chassis.

**Step 3: Replace the chassis and install components**

Remove the impaired chassis, install the replacement chassis, and reinstall all components.

**Steps**

1. Remove the impaired chassis:

- Remove the screws from the chassis mount points.
- Using two people or a lift, 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. Using two people or a lift, 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.

3. Beginning with the bottom controller module, install the controller modules in the replacement chassis:

- a. Align the end of the controller module with the opening in the chassis, and then gently push the controller all the way into the chassis.
- b. Rotate the locking latches upward into the locked position.
- c. If you have not already done so, reinstall the cable management device and recable the controller.

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

Make sure that the cables are connected referencing the cable labels.

4. Reinstall the drives into their corresponding drive bays in the front of the chassis.

5. Install all four of the PSUs:

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

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



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

6. Reconnect the PSU power cables to all four of the PSUs.

- a. Secure the 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 chassis and reinstalling the components, [complete the chassis replacement](#).

### Complete the chassis replacement - ASA A70 and ASA A90

Reboot the controllers, verify system health, and return the failed part to NetApp to complete the final step in the ASA A70 and ASA A90 chassis replacement procedure.

#### Step 1: Boot the controllers and perform giveback

Boot ONTAP on the controllers and perform controller giveback to return storage ownership.

#### Steps

1. Check the console output:
  - a. If the controller boots to the LOADER prompt, reboot the controller with the `boot_ontap` command.
  - b. If the console displays `waiting for giveback` after the reboot, log into the partner controller and check that the replaced controller is ready for giveback with the `storage failover show` command.
2. Perform the giveback:
  - a. Connect the console cable to the partner controller.
  - b. Give back the controller with the `storage failover giveback -fromnode local` command.

### Step 2: Verify storage system health

After the controller giveback completes, verify system health using [Active IQ Config Advisor](#). Address any issues found.

### 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 - ASA A70 and ASA A90

Get started with replacing the controller in your ASA A70 or ASA A90 storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.

1

#### [Review controller replacement requirements](#)

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

2

#### [Shut down the impaired controller](#)

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

3

#### [Replace the controller](#)

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

4

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

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

5

## Recable and give back the controller

Recable the controller and transfer the ownership of storage resources back to the replacement controller.

6

## Complete controller replacement

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

### Requirements to replace the controller - ASA A70 and ASA A90

Before replacing the controller in your ASA A70 or ASA A90 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 module.

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

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

### Shut down the impaired controller - ASA A70 and ASA A90

Shut down the controller in your ASA A70 or ASA A90 storage system to prevent data loss and ensure system stability when replacing the controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take

over the controller so that the healthy controller continues to serve data from the impaired controller storage.

## About this task

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

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

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

## Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

## Replace the controller - ASA A70 and ASA A90

Replace the controller in your ASA A70 or ASA A90 system when a hardware failure requires it. This process involves removing the impaired controller, moving the

components to the replacement controller, installing the replacement controller, and rebooting the system.

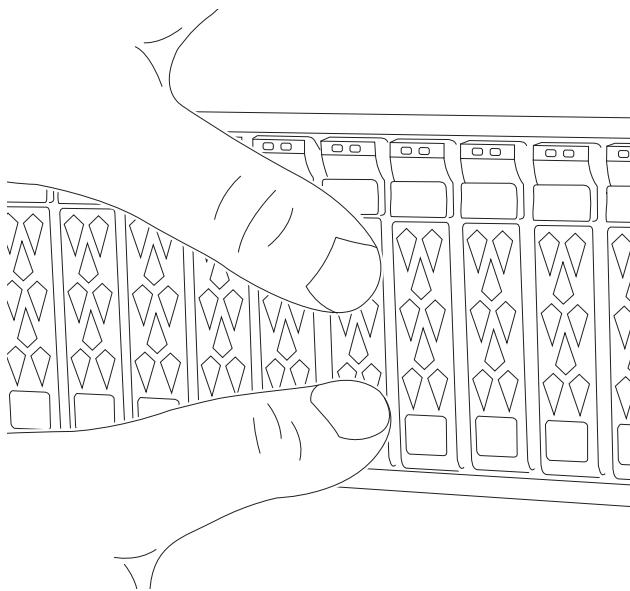
#### Step 1: Remove the controller module

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

#### Steps

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

[Video - Confirm drive seating](#)



3. Check the controller drives based on the system status:

- a. On the healthy controller, check if any active RAID group is in a degraded state, failed state, or both:

```
storage aggregate show -raidstatus !*normal*
```

- If the command returns There are no entries matching your query. continue to [go to the next sub-step to check for missing drives](#).
- If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message
'<message_name>'
```

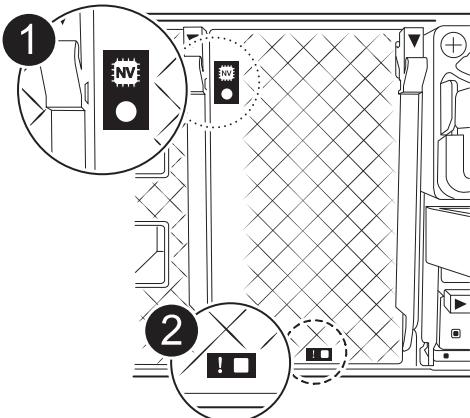
- b. Check for missing drive issues for both the file system or spare drives:

```
event log show -severity * -node * -message-name *disk.missing*
```

- If the command returns There are no entries matching your query. continue to [go to the next step](#).
- If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message  
'<message_name>'
```

4. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. 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.

5. If you are not already grounded, properly ground yourself.

6. Unplug the controller module power supply cables from the controller module power supplies (PSU).



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

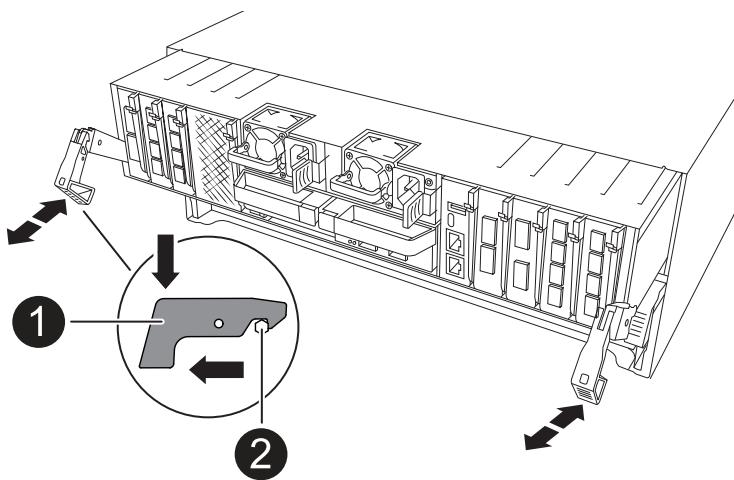
7. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

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

8. Remove the cable management device from the controller module.

9. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.



1	Locking latch
2	Locking pin

10. 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 2: Move the power supplies

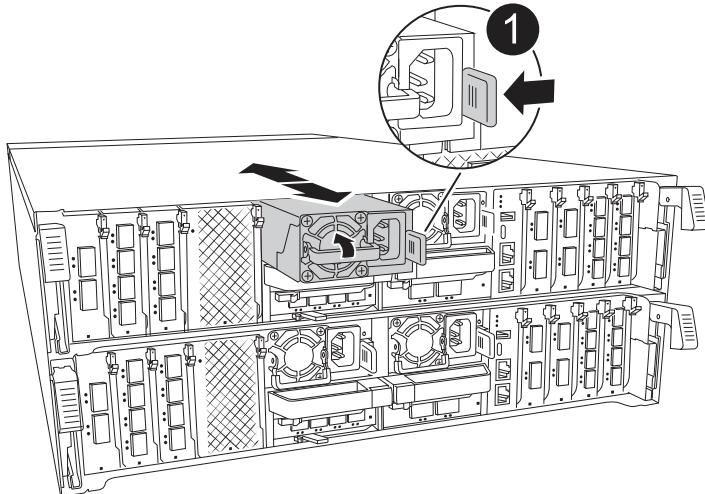
Move the power supplies to the replacement controller.

##### Steps

1. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.



The power supply 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
2	Power supply

2. Move the power supply to the new controller module, and then install it.
3. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

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



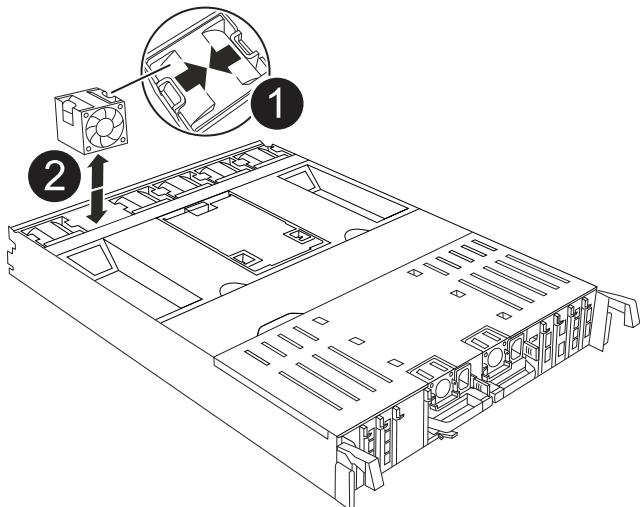
To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

### Step 3: Move the fans

Move the fans modules to the replacement controller module.

#### Steps

1. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.



1	Fan locking tabs
2	Fan module

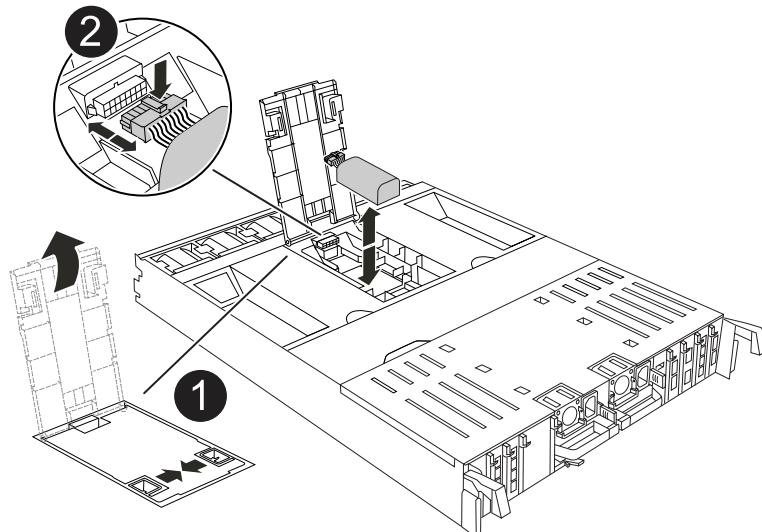
2. Move the fan module to the replacement controller module, and then install it by aligning its edges with the opening in the controller module and sliding it in until the locking latches click into place.
3. Repeat these steps for the remaining fan modules.

#### Step 4: Move the NV battery

Move the NV battery to the replacement controller module.

##### Steps

1. Open the air duct cover in the middle of the controller module and locate the NV battery.



1	NV battery air duct
---	---------------------

2

NV battery pack plug

**Attention:** The NV module LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

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 replacement controller module:
  - 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 NV battery air duct.

#### **Step 5: Move system DIMMs**

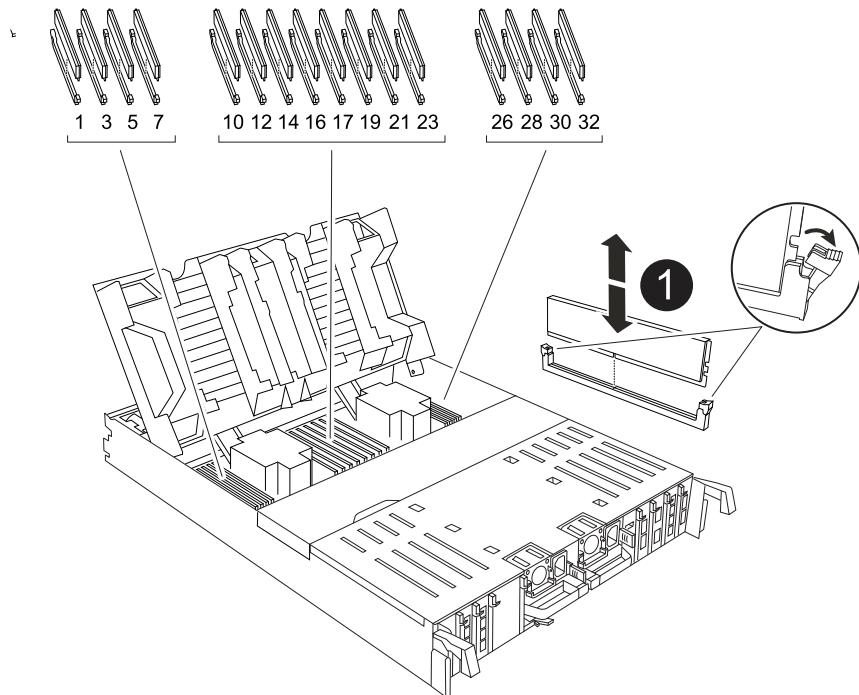
Move the DIMMs to the replacement controller module.

#### **Steps**

1. Open the controller air duct on the top of the controller.
  - a. Insert your fingers in the recesses at the far ends of the air duct.
  - b. Lift the air duct and rotate it upward as far as it will go.
2. Locate the system DIMMs on the motherboard, using the DIMM map on top of the air duct.

The DIMM locations, by model, are listed in the following table:

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



1

System DIMM

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



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

5. Locate the slot on the replacement controller module where you are installing the DIMM.
6. 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. Repeat these steps for the remaining DIMMs.
9. Close the controller air duct.

#### Step 6: Move the I/O modules

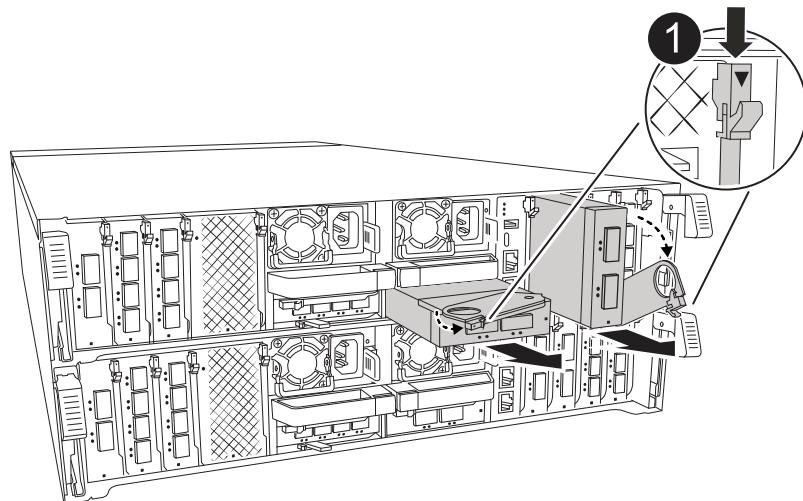
Move the I/O modules to the replacement controller module.

#### Steps

1. Unplug any cabling on the target I/O module.

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

2. Rotate the cable management arm down by pulling the buttons on the inside of the cable management arm and rotating it down.



1

I/O module cam lever

3. Remove the I/O modules from the controller module:

- a. Depress the target I/O module cam latch button.
- b. Rotate the cam latch down as far as it will go. For horizontal modules, rotate the cam 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.

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

- d. Install the replacement I/O module into the replacement controller module by gently sliding the I/O module into the slot until the I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.

4. Repeat these steps to move the remaining I/O modules, except the modules in slots 6 and 7, to the replacement controller module.



To move the I/O modules from slots 6 and 7, you must move the carrier containing these I/O modules from the impaired controller module to the replacement controller module.

5. Move the carrier containing the I/O modules in slots 6 and 7 to the replacement controller module:

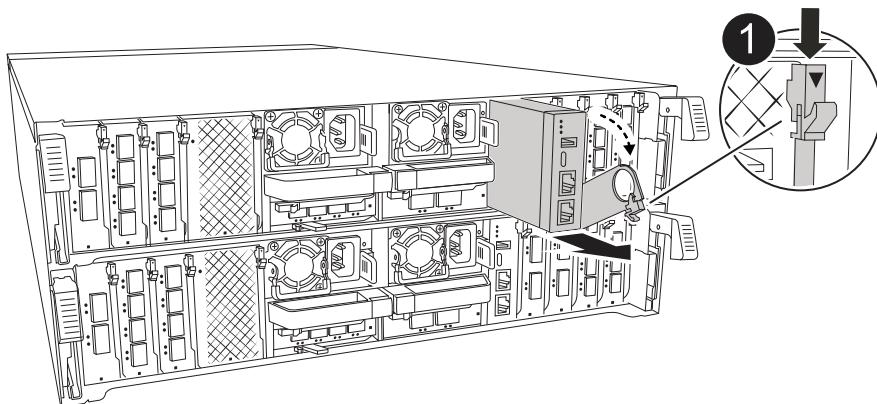
- a. Press the button on the right-most handle on the carrier handle.  
..Slide the carrier out of the impaired controller module insert it into the replacement controller module in the same position it was in the impaired controller module.
- b. Gently push the carrier all the way into the replacement controller module until it locks into place.

#### Step 7: Move the System Management module

Move the System Management module to the replacement controller module.

## Steps

1. Remove the System Management module from the impaired controller module:



**1**

System Management module cam latch

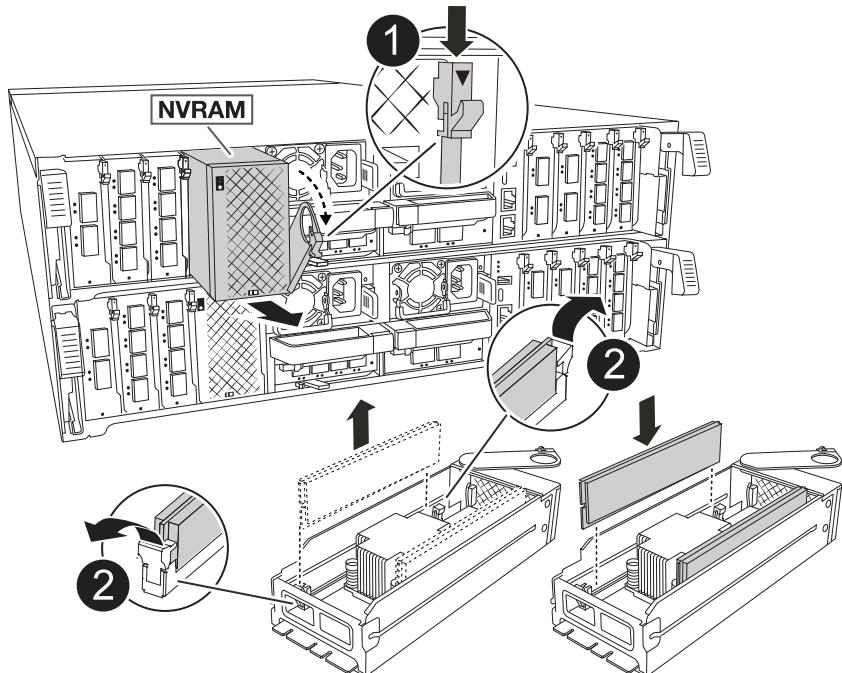
- a. Depress the system management cam button.  
b. Rotate the cam lever all the way down.  
c. Loop your finger into the cam lever and pull the module straight out of the system.
2. Install the system management module into the replacement controller module in the same slot that it was in on the impaired controller module:
  - a. Align the edges of the 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.

### Step 8: Move the NVRAM module

Move the NVRAM module to the replacement controller module.

## Steps

1. Remove the NVRAM module from the impaired controller module:



1	Cam locking button
2	DIMM locking tab

- Depress the cam latch button.

The cam button moves away from the chassis.

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

## 2. Install the NVRAM module into slot 4/5 in the replacement controller module:

- Align the module with the edges of the chassis opening in slot 4/5.
- Gently slide the module into the slot all the way, and then push the cam latch all the way up to lock the module in place.

### Step 9: Install the controller module

Reinstall the controller module and reboot it.

#### Steps

- 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.

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



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

3. Reinstall the cable management arm, if removed, but do not reconnect any cables to the replacement controller.
4. Plug the console cable into the console port of the replacement controller module and reconnect it to the laptop so that it receives console messages when it reboots.
5. Complete the reinstallation of the controller module:
  - a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latches rise when the controller module is fully seated.



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

- a. Rotate the locking latches upward into the locked position.
- b. Plug in the power supplies. The controller boots to the LOADER prompt as soon as power is restored.

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



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

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

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

9. Plug the power cords into the power supplies.



If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

## What's next?

After you've replaced the impaired ASA A70 or ASA A90 controller, you need to [restore the system configuration](#).

## Restore and verify the system configuration - ASA A70 and ASA A90

Verify that the controller's HA configuration is active and functioning correctly in your ASA A70 or ASA A90 storage system, and confirm that the system's adapters list all the paths to the disks.

Verify the HA configuration settings and disk list before returning the system to service.

## Steps

1. Boot to maintenance mode:

```
boot_ontap maint
```

2. Enter **y** when you see *Continue with boot?*.

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

3. Enter **sysconfig -v** and capture the display contents.



If you see *PERSONALITY MISMATCH* contact customer support.

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

5. Verify that all components display the same HA state:

```
ha-config show
```

The HA state should be the same for all components.

6. If the displayed system state of the controller module does not match your system configuration, set the HA state for the controller module:

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

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

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

7. Confirm that the setting has changed:

```
ha-config show
```

8. Verify that the adapter lists the paths to all disks: **storage show disk -p**

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

9. Exit Maintenance mode: **halt**

### Give back the controller - ASA A70 and ASA A90

Return control of storage resources to the replacement controller so your ASA A70 or ASA A90 system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption or Onboard Key Manager (OKM) encryption.

## No encryption

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

### Steps

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

## Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

### Steps

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



You are prompted twice for the passphrase.

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



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

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

The output should show no results when when properly synchronized.

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

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

## Complete controller replacement - ASA A70 and ASA A90

To complete the controller replacement for your ASA A70 or ASA A90 system, first restore the NetApp Storage Encryption configuration (if necessary). Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, return the failed part to NetApp.

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

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

#### Steps

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

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

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

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

2. Check the health of your cluster. See the [How to perform a cluster health check with a script in ONTAP](#) KB article.

3. If automatic giveback was disabled, reenable it:

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

### Step 2: Return the failed part to NetApp

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

## Replace a DIMM - ASA A70 and ASA A90

Replace a DIMM in your ASA A70 or ASA A90 system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting ONTAP. The replacement process involves shutting down the impaired

controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

#### Before you begin

- All other components in the system must be functioning properly; if not, you must contact technical support.
- You must replace the failed component with a replacement FRU component you received from your provider.

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

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, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

#### About this task

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

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

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

#### Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

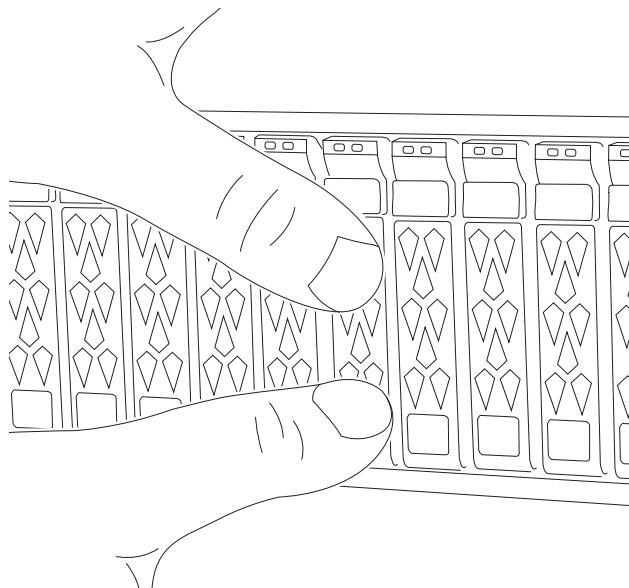
If the impaired controller is displaying...	Then...
Waiting for giveback...	Press Ctrl-C, and then respond <b>y</b> when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: <code>storage failover takeover -ofnode impaired_node_name -halt true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

## Step 2: Remove the controller module

### Steps

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

[Video - Confirm drive seating](#)



3. Check the controller drives based on the system status:

- a. On the healthy controller, check if any active RAID group is in a degraded state, failed state, or both:

```
storage aggregate show -raidstatus !*normal*
```

- If the command returns There are no entries matching your query. continue to [go to the next sub-step to check for missing drives](#).
- If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message  
'<message_name>'
```

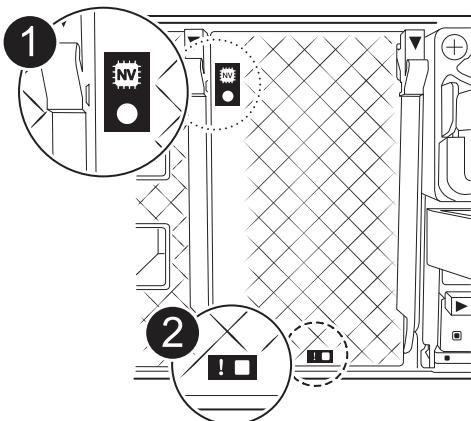
b. Check for missing drive issues for both the file system or spare drives:

```
event log show -severity * -node * -message-name *disk.missing*
```

- If the command returns There are no entries matching your query. continue to [go to the next step](#).
- If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message  
'<message_name>'
```

4. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. 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.

5. If you are not already grounded, properly ground yourself.

6. Unplug the controller module power supply cables from the controller module power supplies (PSU).



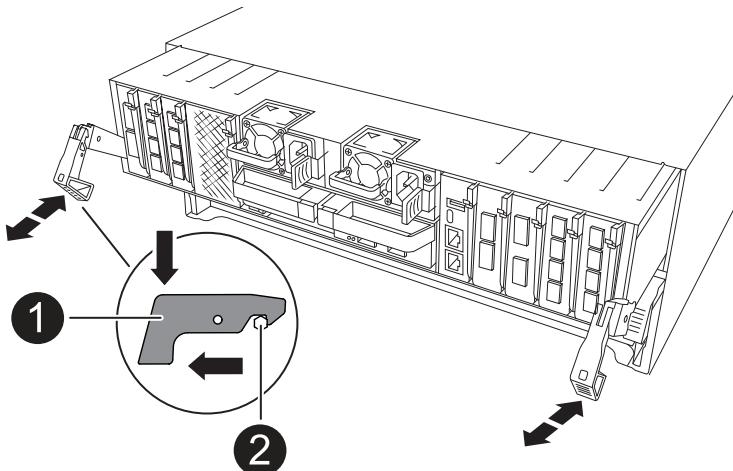
If your system has DC power, disconnect the power block from the PSUs.

7. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

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

8. Remove the cable management device from the controller module.
9. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.



1	Locking latch
2	Locking pin

10. 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 3: Replace a DIMM

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

#### Steps

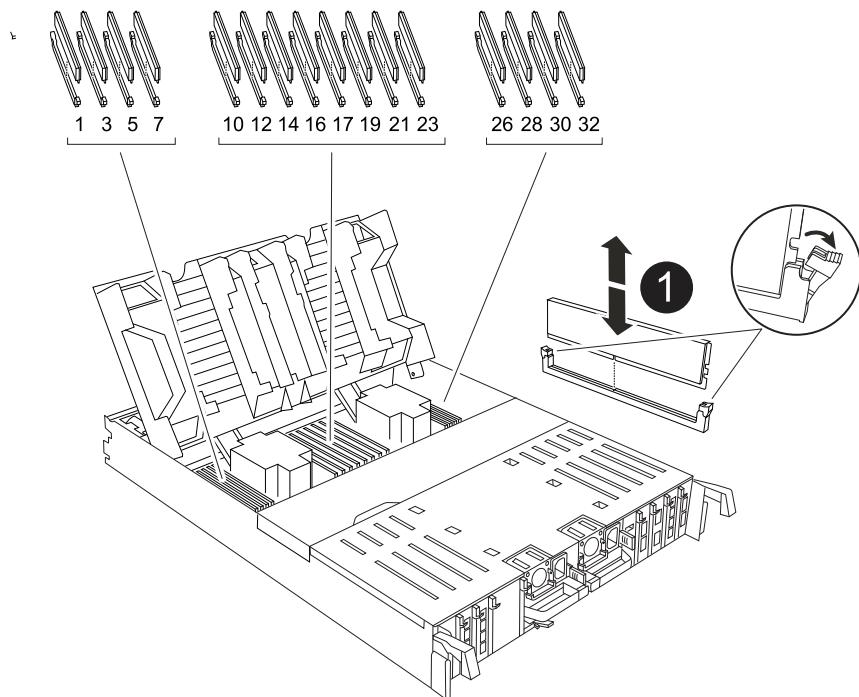
1. If you are not already grounded, properly ground yourself.
2. Open the controller air duct on the top of the controller.
  - a. Insert your fingers in the recesses at the far ends of the air duct.
  - b. Lift the air duct and rotate it upward as far as it will go.
3. Locate the DIMMs on your controller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

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.



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: Reinstall the controller module

Reinstall the controller module and reboot it.

##### Steps

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

It must lie flush against the controller module sheet metal.

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



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

3. Recable the storage system, as needed.

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

Make sure that the console cable is connected to the repaired controller module so that it receives console messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. Complete the reinstallation of the controller module:

- a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latches rise when the controller module is fully seated.

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

- b. Rotate the locking latches upward into the locked position.

5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

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

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

7. If automatic giveback was disabled, reenable it:

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

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

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

## 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 an SSD drive - ASA A70 and ASA A90

Replace a drive in your ASA A70 or ASA A90 system when a drive fails or requires an upgrade. This process involves identifying the faulty drive, safely removing it, and installing a new drive to ensure continued data access and system performance.

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

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

### Before you begin

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

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

It can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

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

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

### About this task

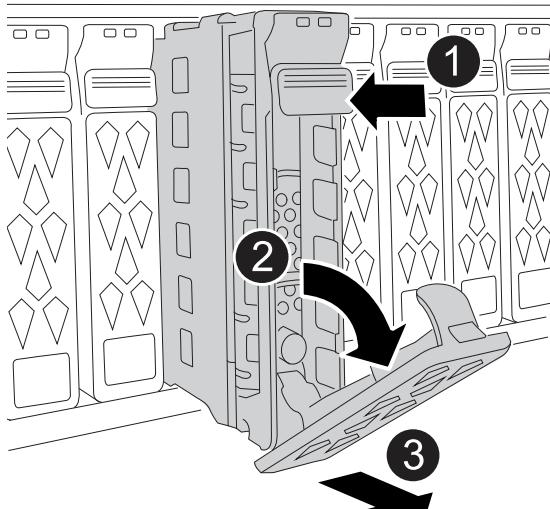
Drive firmware is automatically updated (nondisruptively) on new drives that have non-current firmware versions.

When replacing several disk drives, you must wait 70 seconds between the removal of each failed disk drive and the insertion of the replacement disk drive to allow the storage system to recognize the existence of each new disk.

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

### Steps

1. Properly ground yourself.
2. Remove the bezel from the front of the storage system.
3. Physically identify the failed drive.
  - When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the shelf operator display panel and the failed drive illuminate.
  - The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:



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

5. Wait a minimum of 70 seconds before inserting the replacement drive.

This allows the system to recognize that a drive was removed.

6. Insert the replacement drive:

- With the cam handle in the open position, use both hands to insert the replacement drive.
- Gently push until the drive stops.
- Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

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

7. Verify that the drive's activity (green) LED is illuminated.

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

- If you are replacing another drive, repeat Steps 3 through 7.
- Reinstall the bezel on the front of the storage system.
- Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Contact technical support [NetApp Support](#) if you need the RMA number or additional help with the replacement procedure.

## Replace a fan module - ASA A70 and ASA A90

Replace a fan module in your ASA A70 or ASA A90 system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

### 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, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

#### About this task

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

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

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

#### Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

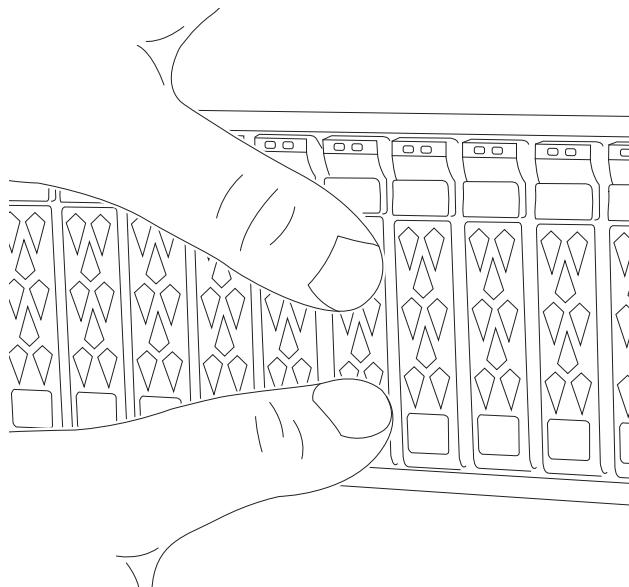
If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <b>y</b> when prompted.
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover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

## Step 2: Remove the controller module

### Steps

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

[Video - Confirm drive seating](#)



3. Check the controller drives based on the system status:

- a. On the healthy controller, check if any active RAID group is in a degraded state, failed state, or both:

```
storage aggregate show -raidstatus !*normal*
```

- If the command returns There are no entries matching your query. continue to [go to the next sub-step to check for missing drives](#).

- If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message
'<message_name>'
```

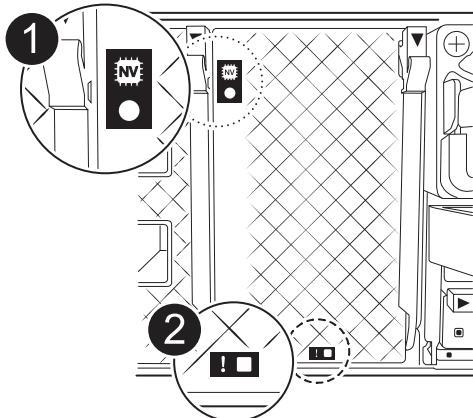
- Check for missing drive issues for both the file system or spare drives:

```
event log show -severity * -node * -message-name *disk.missing*
```

- If the command returns There are no entries matching your query. continue to [go to the next step](#).
- If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message
'<message_name>'
```

- Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. 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.

- If you are not already grounded, properly ground yourself.
- Unplug the controller module power supply cables from the controller module power supplies (PSU).



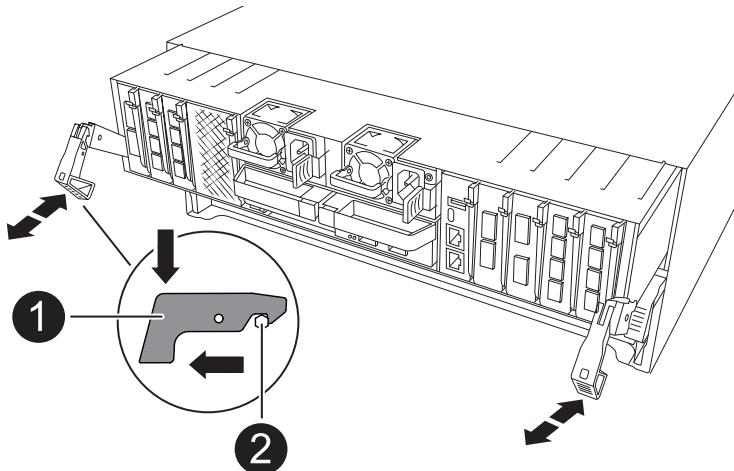
If your system has DC power, disconnect the power block from the PSUs.

7. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

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

8. Remove the cable management device from the controller module.
9. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.



1	Locking latch
2	Locking pin

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

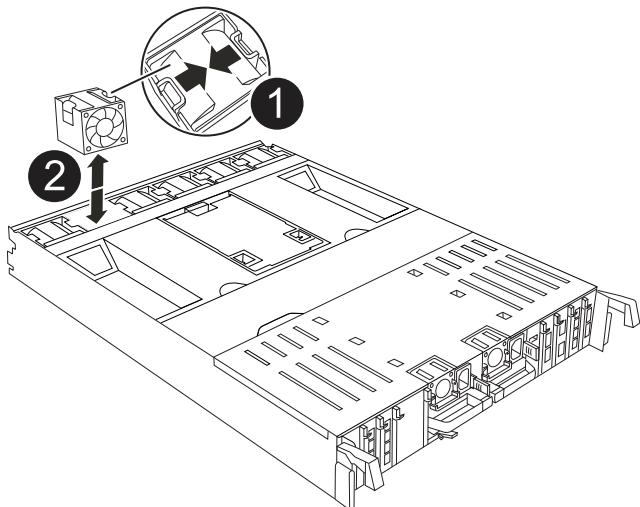
Make sure that you support the bottom of the controller module as you slide it out of the chassis.

### Step 3: Replace the fan

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

#### Steps

1. Identify the fan module that you must replace by checking the console error messages.
2. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.



1	Fan locking tabs
2	Fan module

3. Align the edges of the replacement fan module with the opening in the controller module, and then slide the replacement fan module into the controller module until the locking latches click into place.

#### Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

##### Steps

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

It must lie flush against the controller module sheet metal.

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



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

3. Recable the storage system, as needed.

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

Make sure that the console cable is connected to the repaired controller module so that it receives console messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. Complete the reinstallation of the controller module:

- a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latches rise when the controller module is fully seated.

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

- b. Rotate the locking latches upward into the locked position.
5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

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

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

7. If automatic giveback was disabled, reenable it:

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

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

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

## 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 NVRAM - ASA A70 and ASA A90

Replace the NVRAM in your ASA A70 or ASA A90 system when the non-volatile memory becomes faulty or requires an upgrade. The replacement process involves shutting down the impaired controller, completely removing power by pulling the controller out a few inches, replacing the NVRAM module or the NVRAM DIMM, reassigning the disks, and returning the failed part to NetApp.

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

### Before you begin

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

### Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

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

### About this task

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

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

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

## Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

## Step 2: Replace the NVRAM module or NVRAM DIMM

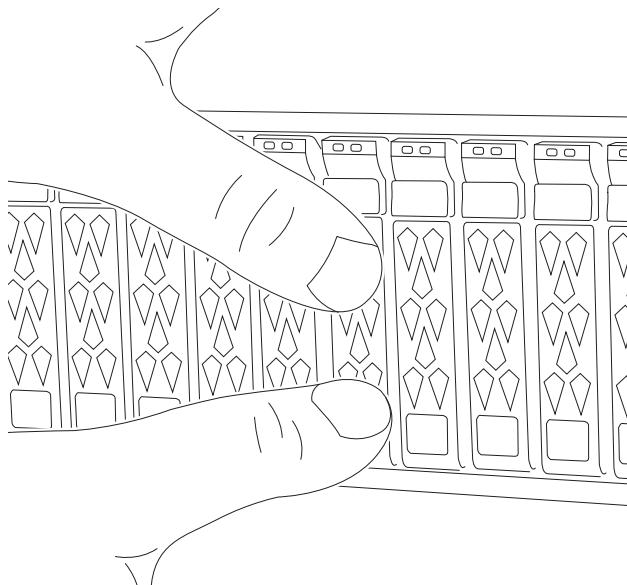
Replace the NVRAM module or NVRAM DIMMs using the appropriate option.

## Option 1: Replace the NVRAM module

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

### Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

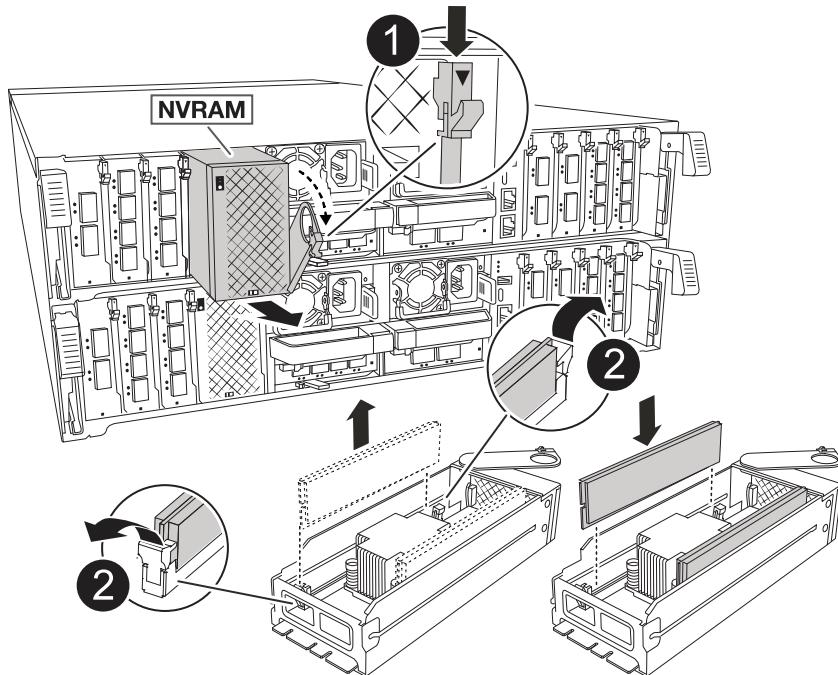


2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. If you are not already grounded, properly ground yourself.
4. Unplug the controller's PSUs.



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

5. Press down on both of the locking latches on the controller, rotate both latches downward at the same time and pull the controller out about 3 or 4 inches.
6. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
7. Remove the target NVRAM module from the chassis:
  - a. Depress the cam latch button.  
The cam button moves away from the chassis.
  - b. Rotate the cam latch as far as it will go.
  - c. Remove the impaired NVRAM module from the chassis by hooking your finger into the cam lever opening and pulling the module out of the chassis.



1	Cam locking button
2	DIMM locking tabs

8. Set the NVRAM module on a stable surface.
9. Remove the DIMMs, one at a time, from the impaired NVRAM module and install them in the replacement NVRAM module.
10. Install the replacement NVRAM module into the chassis:
  - a. Align the module with the edges of the chassis opening in slot 4/5.
  - b. Gently slide the module into the slot all the way, and then push the cam latch all the way up to lock the module in place.
11. Reinstall the controller module. The controller reboots as soon as it is fully inserted.
  - a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latches rise when the controller module is fully seated.



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

- b. Rotate the locking latches upward into the locked position.
12. Plug the power cords into the power supplies.



If you have DC power supplies, reconnect the power block to the power supplies.

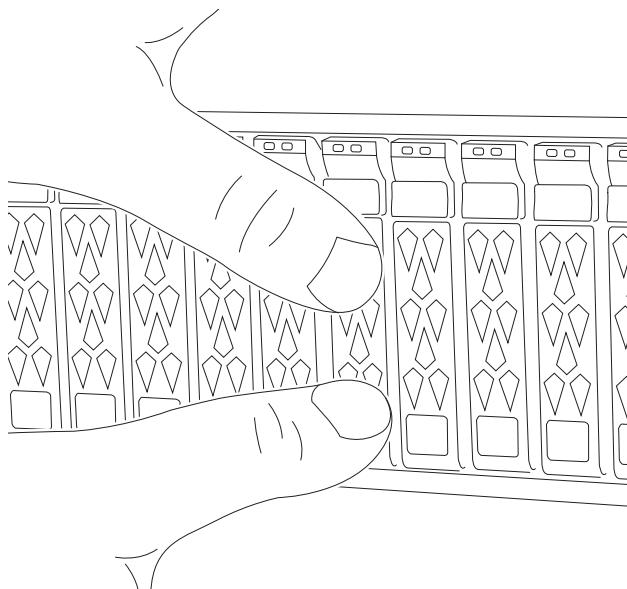
13. Rotate the cable management tray up to the closed position.

#### Option 2: Replace the NVRAM DIMM

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

## Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.



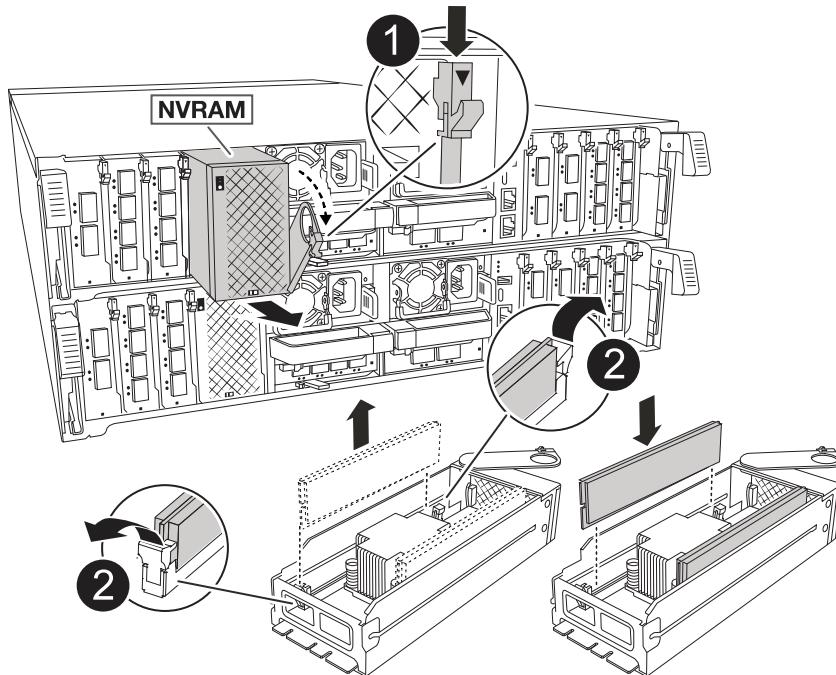
2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Unplug the controller's PSUs.



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

4. Press down on both of the locking latches on the controller, rotate both latches downward at the same time and pull the controller out about 3 or 4 inches.
5. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the tray down.
6. Remove the target NVRAM module from the chassis:
  - a. Depress the cam button.

The cam button moves away from the chassis.
  - b. Rotate the cam latch 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	Cam locking button
2	DIMM locking tabs

7. Set the NVRAM module on a stable surface.
8. 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.

9. Remove the DIMM by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.
10. 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.
11. Install the NVRAM module into the chassis:
  - 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.
12. Reinstall the controller module. The controller reboots as soon as it is fully inserted.
  - Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latches rise when the controller module is fully seated.



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

- Rotate the locking latches upward into the locked position.

13. Plug the power cords into the power supplies.



If you have DC power supplies, reconnect the power block to the power supplies.

14. Rotate the cable management tray up to the closed position.

### Step 3: Verify controller state

You must confirm the controller state of the controllers connected to the disk pool when you boot the controller.

#### Steps

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

In the command output, you should see a message indicates the state of the controllers.

Node	Partner	Possible State	Description
-----	-----	-----	-----
-----	-----	-----	-----
<nodename>			
	<nodename>- P2-3-178	true	Connected to <nodename>-P2-3-178.
to			Waiting for cluster applications
			come online on the local node.
AFF-A90-NBC-P2-3-178			
	<nodename>- P2-3-177	true	Connected to <nodename>-P2-3-177,
2 entries were displayed.			Partial giveback

4. Give back the controller:

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

The controller connects back its storage pool and completes booting.

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



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

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

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

## 5. Verify all disks are displayed: `storage disk show`

```
 ::> storage disk show
      Usable          Disk      Container  Container
Disk      Size Shelf Bay Type      Type      Name
-----
1.0.0      3.49TB   0   0 SSD-NVM aggregate  pod_NVME_SSD_1
1.0.1      3.49TB   0   1 SSD-NVM aggregate  pod_NVME_SSD_1
1.0.2      3.49TB   0   2 SSD-NVM aggregate  pod_NVME_SSD_1
1.0.3      3.49TB   0   3 SSD-NVM aggregate  pod_NVME_SSD_1
1.0.4      3.49TB   0   4 SSD-NVM aggregate  pod_NVME_SSD_1
[...]
48 entries were displayed.
```

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

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

## Replace the NV battery - ASA A70 and ASA A90

Replace the NV battery in your ASA A70 or ASA A90 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, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

### About this task

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

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

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

impaired controller; see [Synchronize a node with the cluster](#).

## Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

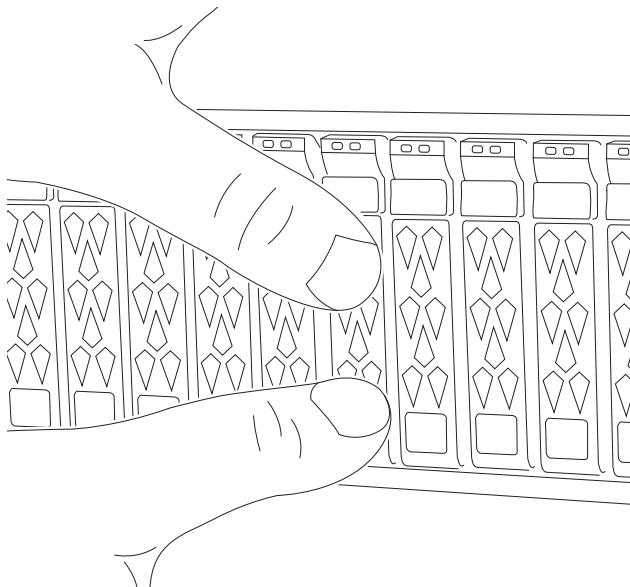
If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller:  <i>storage failover takeover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

## Step 2: Remove the controller module

### Steps

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

[Video - Confirm drive seating](#)



3. Check the controller drives based on the system status:

a. On the healthy controller, check if any active RAID group is in a degraded state, failed state, or both:

```
storage aggregate show -raidstatus !*normal*
```

- If the command returns There are no entries matching your query. continue to [go to the next sub-step to check for missing drives](#).
- If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message  
'<message_name>'
```

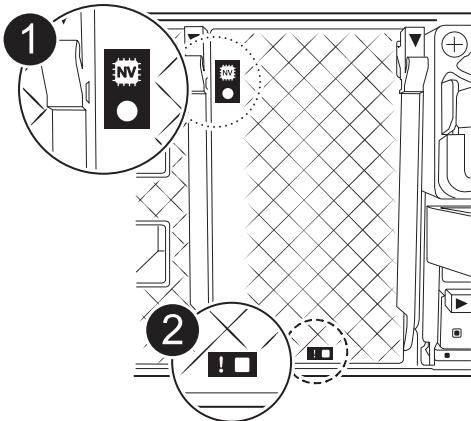
b. Check for missing drive issues for both the file system or spare drives:

```
event log show -severity * -node * -message-name *disk.missing*
```

- If the command returns There are no entries matching your query. continue to [go to the next step](#).
- If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message  
'<message_name>'
```

4. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. 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.

5. If you are not already grounded, properly ground yourself.
6. Unplug the controller module power supply cables from the controller module power supplies (PSU).



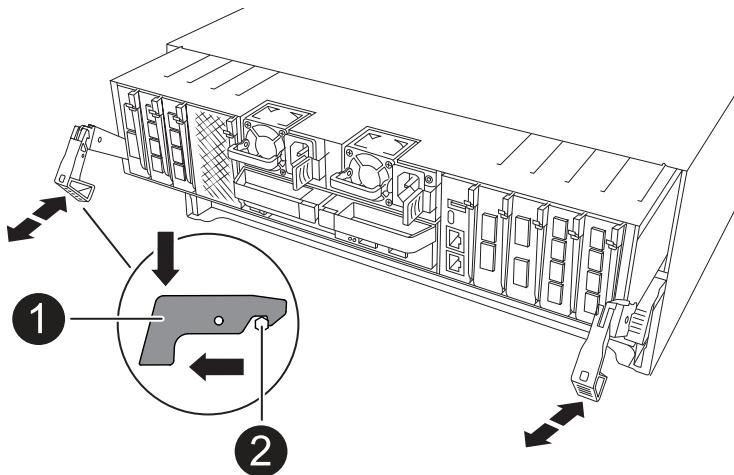
If your system has DC power, disconnect the power block from the PSUs.

7. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

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

8. Remove the cable management device from the controller module.
9. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.



1	Locking latch
2	Locking pin

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

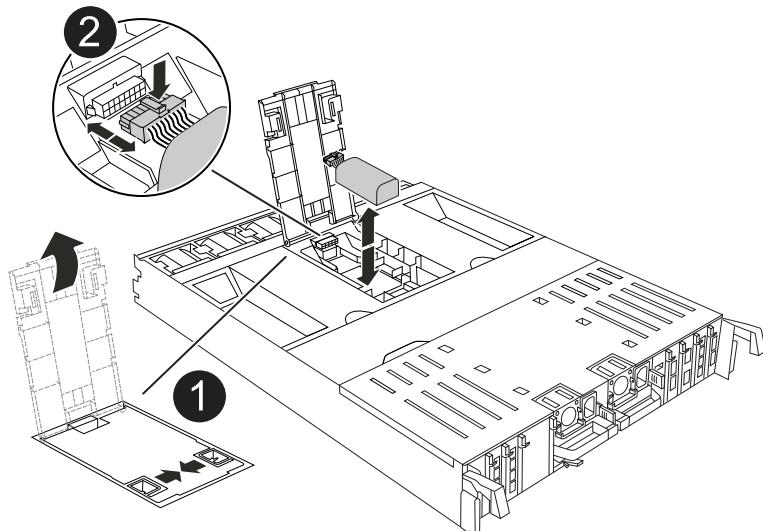
Make sure that you support the bottom of the controller module as you slide it out of the chassis.

### Step 3: Replace the NV battery

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

#### Steps

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



1	NV battery air duct cover
---	---------------------------

2

NV battery plug

2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module, and then set it aside.
5. Remove the replacement battery from its package.
6. Install the replacement battery pack into the controller:
  - a. Plug the battery plug into the riser socket and make sure that the plug locks into place.
  - b. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
7. Close the NV air duct cover.

Make sure that the plug locks into the socket.

#### Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

##### Steps

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

It must lie flush against the controller module sheet metal.

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



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

3. Recable the storage system, as needed.

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

Make sure that the console cable is connected to the repaired controller module so that it receives console messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. Complete the reinstallation of the controller module:

- a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latches rise when the controller module is fully seated.

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

- b. Rotate the locking latches upward into the locked position.

5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

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

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

7. If automatic giveback was disabled, reenable it:

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

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

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

## Step 5: Return the failed part to NetApp

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

## I/O module

### Overview of add and replace an I/O module - ASA A70 and ASA A90

The ASA A70 and ASA A90 systems offer 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 ASA A70 or ASA A90 storage system with the same type of I/O module, or with a different kind of I/O module. You can also add an I/O module into a system with empty slots.

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

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

- [Hot-swap an I/O module](#)

You can hot-swap certain I/O modules for an equivalent I/O module to restore the storage system to its optimal operating state. Hot-swapping is done without having to perform a manual takeover.

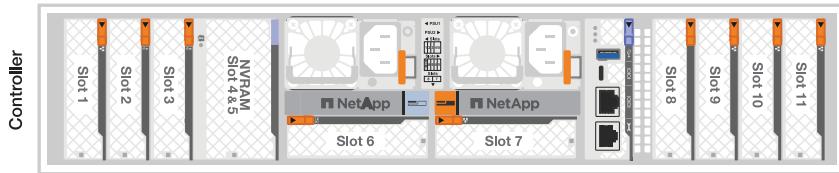
To use this procedure, your storage system must be running ONTAP 9.18.1.

- [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 ASA A70 and ASA A90 controllers are numbered 1 through 11, as shown in the following illustration.



## Add an I/O module - ASA A70 and ASA A90

Add an I/O module to your ASA A70 or ASA A90 system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your ASA A70 and ASA A90 storage system when there are empty slots available or when all slots are fully populated.

### About this task

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

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

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

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

Shut down or take over 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.

### About this task

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

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

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

### Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

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

#### Before you begin

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

## Add I/O module to an available slot

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

### Steps

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



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

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

bye



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

8. Give back the controller from the partner controller:

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

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

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

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

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

## Add I/O module to a fully-populated system

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

### About this task

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

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See <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.



If you encounter an issue during reboot, see [BURT 1494308 - Environment shutdown might be triggered during I/O module replacement](#)

10. Give back the controller from the partner controller:

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

11. Enable automatic giveback if it was disabled:

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

12. Do one of the following:

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

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

- If you removed a NIC I/O module and installed a storage I/O module, install and cable your NS224 shelves, as described in [Hot-add workflow](#).

13. Repeat these steps for controller B.

## Hot swap an I/O module - ASA A70 and ASA A90

You can hot swap an Ethernet I/O module in your ASA A70 or ASA A90 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 <i>y</i> 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: Replace the failed I/O module**

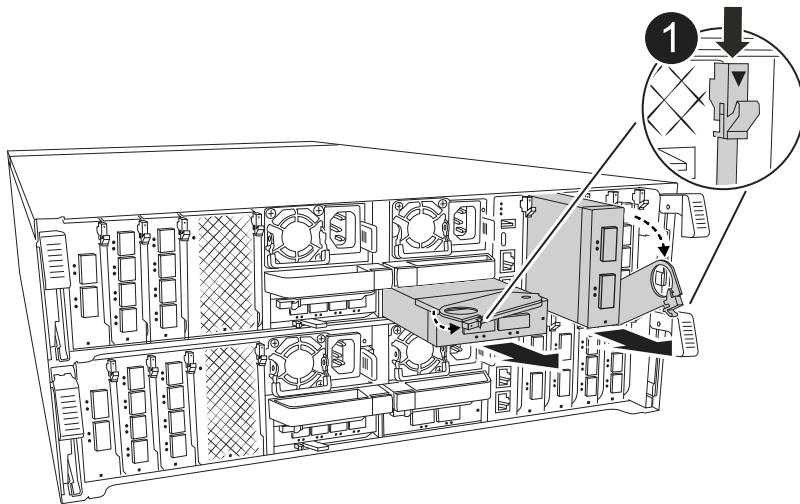
Replace the failed I/O module with an equivalent I/O module.

**Steps**

1. If you are not already grounded, properly ground yourself.
2. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
3. Remove the I/O module from the controller module:



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



<b>1</b>	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 an I/O module - ASA A70 and ASA A90

Replace an I/O module in your ASA A70 or ASA A90 system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

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

#### Before you begin

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

#### Step 1: Shut down the impaired 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, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

#### About this task

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

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

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a

healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

## Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

## Step 2: Replace a failed I/O module

### Steps

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

### Steps

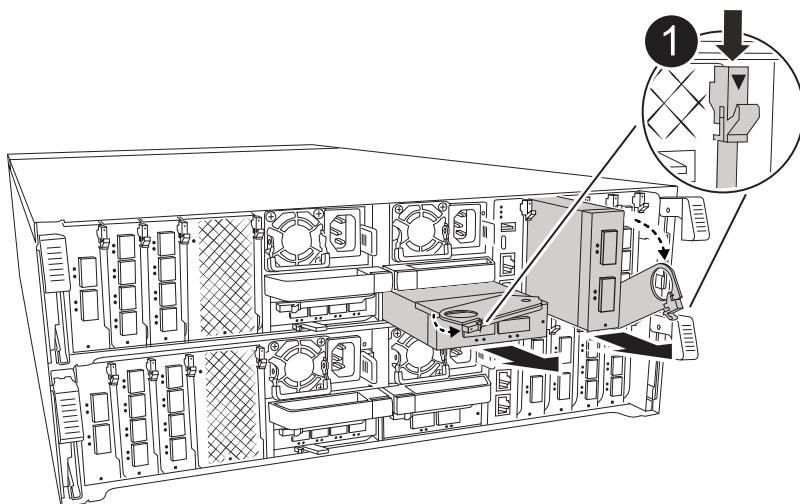
1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.

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

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 I/O module from the controller module:



This 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 do 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.

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 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.
7. Cable the I/O module.
8. Repeat the remove and install steps to replace additional modules for the controller.
9. Rotate the cable management tray into the locked position.

### Step 3: Reboot the controller

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

#### Steps

1. Reboot the controller from the LOADER prompt:

bye



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

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

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

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

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

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

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

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

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

## **Hot-swap a power supply - ASA A70 and ASA A90**

Replace an AC or DC power supply unit (PSU) in your ASA A70 or ASA A90 system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cable, replacing the faulty PSU, and then reconnecting it to the power source.

The power supplies are redundant and hot-swappable. You do not have to shut down the controller to replace a PSU.

### **About this task**

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



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

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

## Option 1: Hot-swap an AC PSU

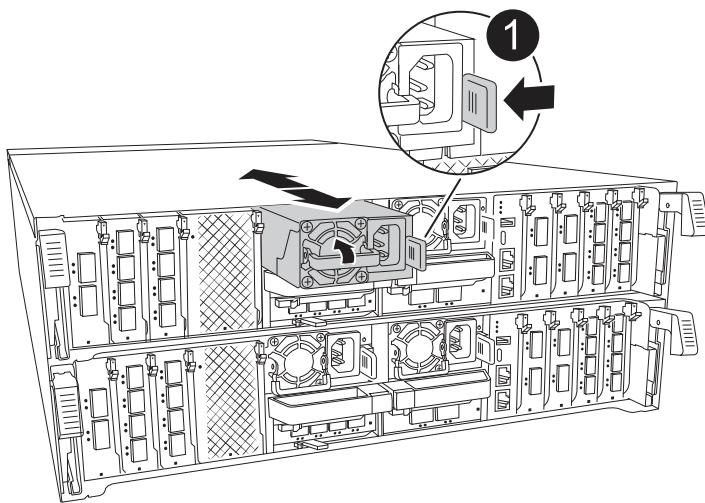
To replace an AC PSU, complete the following steps.

### Steps

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



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



1

Terracotta PSU locking tab

5. Install the replacement PSU in the controller module:

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

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

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

6. Reconnect the PSU cabling:

- a. Reconnect the power cable to the PSU.

- b. Secure the power cable to the PSU using the power cable retainer.

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

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

## Option 2: Hot-swap a DC PSU

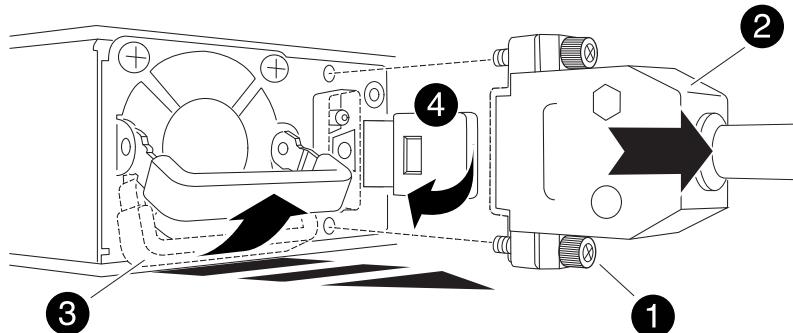
To replace a DC PSU, complete the following steps.

### Steps

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



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



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

5. Install the replacement PSU in the controller module:

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

- b. Gently push the PSU into the controller module until the locking tab clicks into place.

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



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

6. Reconnect the D-SUB DC power cable:

- a. Plug the power cable connector into the PSU.
- b. Secure the power cable to the PSU with the thumbscrews.

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

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

## Replace the real-time clock battery - ASA A70 and ASA A90

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your ASA A70 or ASA A90 system to ensure that services and applications relying on accurate time synchronization remain operational.

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.

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

#### About this task

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

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

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

#### Steps

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

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

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

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

## 2. Disable automatic giveback:

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

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

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

## 3. Take the impaired controller to the LOADER prompt:

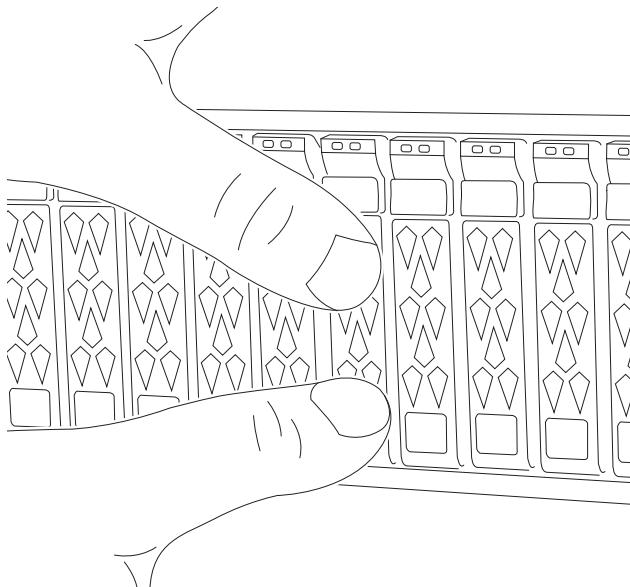
If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller:  <pre>storage failover takeover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

## Step 2: Remove the controller module

### Steps

- If you are not already grounded, properly ground yourself.
- Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

[Video - Confirm drive seating](#)



3. Check the controller drives based on the system status:

a. On the healthy controller, check if any active RAID group is in a degraded state, failed state, or both:

```
storage aggregate show -raidstatus !*normal*
```

- If the command returns There are no entries matching your query. continue to [go to the next sub-step to check for missing drives](#).
- If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message  
'<message_name>'
```

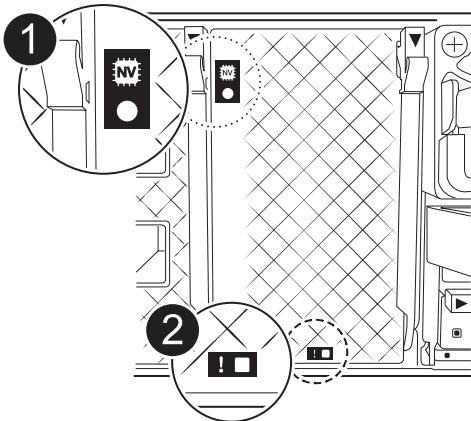
b. Check for missing drive issues for both the file system or spare drives:

```
event log show -severity * -node * -message-name *disk.missing*
```

- If the command returns There are no entries matching your query. continue to [go to the next step](#).
- If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message  
'<message_name>'
```

4. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. 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.

5. If you are not already grounded, properly ground yourself.
6. Unplug the controller module power supply cables from the controller module power supplies (PSU).



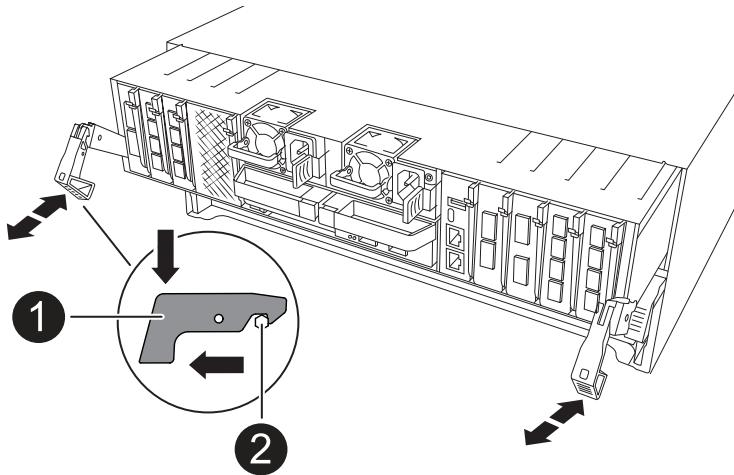
If your system has DC power, disconnect the power block from the PSUs.

7. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

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

8. Remove the cable management device from the controller module.
9. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.



1	Locking latch
2	Locking pin

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

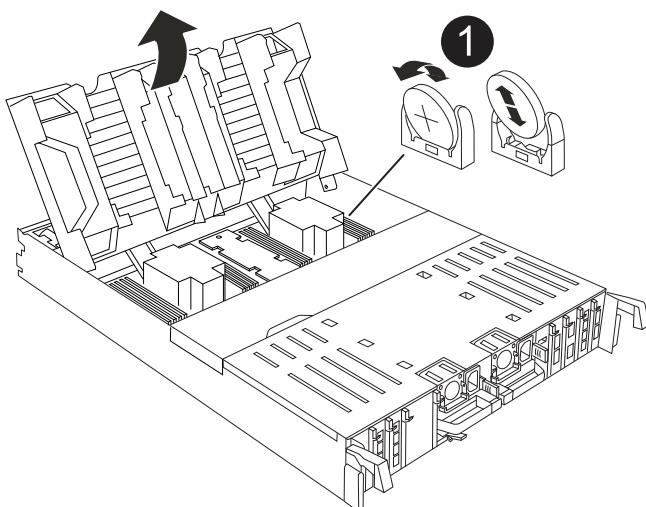
Make sure that you support the bottom of the controller module as you slide it out of the chassis.

### Step 3: Replace the RTC battery

Remove failed RTC battery and install the replacement RTC battery.

#### Steps

1. Open the controller air duct on the top of the controller.
  - a. Insert your fingers in the recesses at the far ends of the air duct.
  - b. Lift the air duct and rotate it upward as far as it will go.
2. Locate the RTC battery under the air duct.



1

RTC battery and housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. 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 reboot it.

##### Steps

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

It must lie flush against the controller module sheet metal.

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



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

3. Recable the storage system, as needed.

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

Make sure that the console cable is connected to the repaired controller module so that it receives console messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. Complete the reinstallation of the controller module:

- a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latches rise when the controller module is fully seated.

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

- b. Rotate the locking latches upward into the locked position.

5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

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

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

7. If automatic giveback was disabled, reenable it:

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

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

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

## Step 5: Reset the time and date on the controller

After you replace the RTC battery, insert the controller, and power on for the first BIOS reset, you will see the following error messages:

RTC date/time error. Reset date/time to default

RTC power failure error

These messages are expected and you can continue with this procedure.

### Steps

1. Check the date and time on the healthy controller with the `cluster date show` command.

If your system stops at the boot menu, select the option for `Reboot node` and respond `y` when prompted, then boot to LOADER by pressing `Ctrl-C`

- a. At the LOADER prompt on the target controller, check the time and date with the `cluster date show` command.
- b. If necessary, modify the date with the `set date mm/dd/yyyy` command.
- c. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.

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.

## 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 the system management module - ASA A70 and ASA A90

Replace the System Management module in your ASA A70 or ASA A90 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.

### Before you begin

- 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.
- All other system components must be working properly.
- The partner controller must be able to take over the impaired controller.
- You must replace the failed component with a replacement FRU component you received from your provider.

## About this task

The System Management module, located at the back of the controller in slot 8, contains onboard components for system management, as well as ports for external management. The target controller must be shut down to replace an impaired System Management module or replace the boot media.

### 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, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

## About this task

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

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

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

## Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

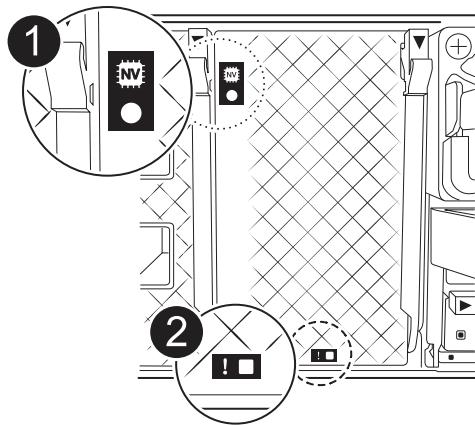
If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

## Step 2: Replace the System Management module

Replace the impaired system management module.

### Steps

1. 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	NVRAM status LED
2	NVRAM attention LED

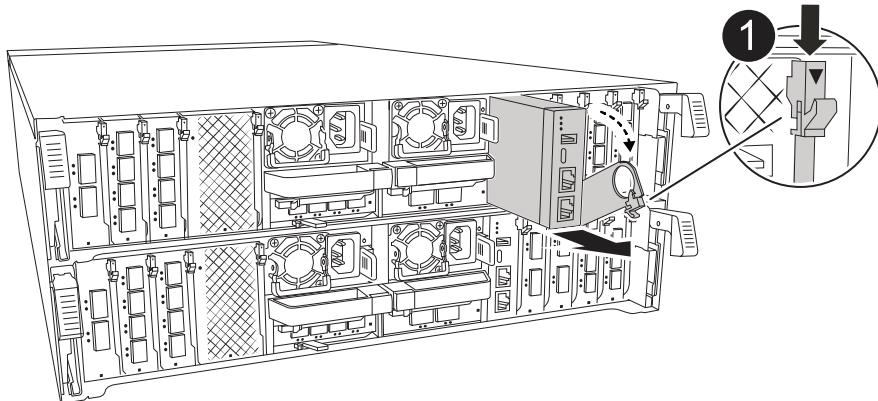
- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Unplug the controller's PSUs.



If your system has DC power, disconnect the power block from the PSUs.

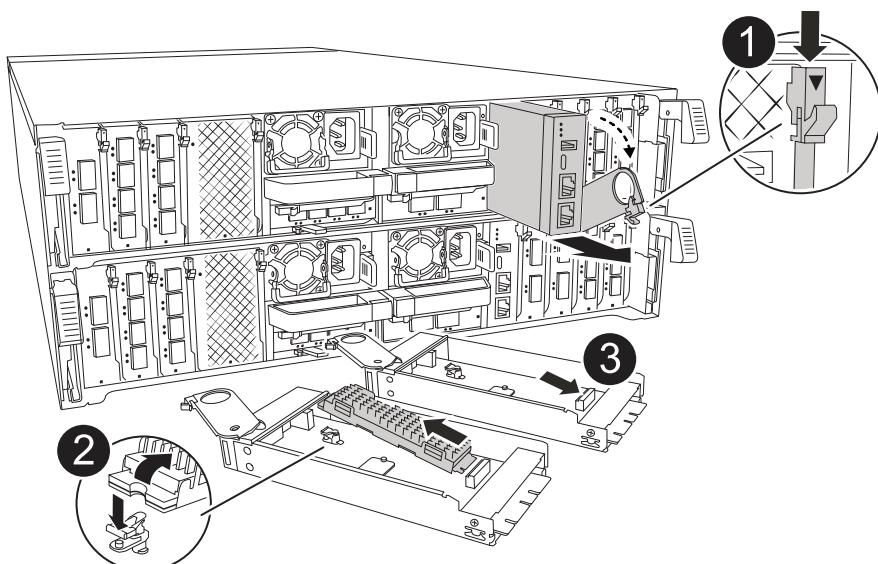
4. Press down on both of the locking latches on the controller, rotate both latches downward at the same time and pull the controller out about 3 or 4 inches.
5. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
6. Remove any cables connected to the System Management module. Make sure that label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.



1

System Management module cam latch

7. Remove the System Management module:
  - a. Depress the system management cam button.  
The cam lever moves away from the chassis.
  - b. Rotate the cam lever all the way down.
  - c. Loop your finger into the cam lever and pull the module straight out of the system.
  - d. Place the System Management module on an anti-static mat, so that the boot media is accessible.
8. 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 locking button.  
The boot media rotates slightly upward.
- b. Rotate the boot media up, slide it out of the socket.
- c. Install the boot media in the replacement System Management module:
  - i. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
  - ii. Rotate the boot media down toward until it engages the locking button. Depress the blue locking if necessary.
9. Install the system management module:
  - 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.
10. Recable the System Management module.
11. Reinstall the controller module. The controller reboots as soon as it is fully inserted.
  - a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.  
The locking latches rise when the controller module is fully seated.



Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

- b. Rotate the locking latches upward into the locked position.

12. Plug the power cords into the power supplies.



If you have DC power supplies, reconnect the power block to the power supplies.

13. Rotate the cable management tray up to the closed position.

### Step 3: Reboot the controller

Reboot the controller module.

#### Steps

1. Enter *bye* at the LOADER prompt.
2. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

### 3. Restore automatic giveback:

```
storage failover modify -node local -auto-giveback true
```

### 4. If an AutoSupport maintenance window was triggered, end it:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

## Step 4: Install licenses and register serial number

You must install new licenses for the node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

### About this task

Until you install license keys, features requiring standard licenses continue to be available to the node. However, if the node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed. Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the node as soon as possible.

### Before you begin

The license keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

 If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in [Post Motherboard Replacement Process to update Licensing on a AFF/FAS system](#). If you are unsure of the initial ONTAP release for your system, see [NetApp Hardware Universe](#) for more information.

### Steps

1. If you need new license keys, obtain replacement license keys on the [NetApp Support Site](#) in the My Support section under Software licenses.



The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
  - a. Check for unused licenses: `license clean-up -unused -simulate`
  - b. If the list looks correct, remove the unused licenses: `license clean-up -unused`
4. Register the system serial number with NetApp Support.
  - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
  - If AutoSupport is not enabled, call [NetApp Support](#) to register the serial number.

## Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

# ASA A20, ASA A30, and ASA A50 systems

## Overview of hardware maintenance - ASA A20, ASA A30, and ASA A50

Maintain the hardware of your ASA A20, ASA A30, or ASA A50 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 ASA A20, ASA A30, and ASA A50 systems have already been deployed as a storage nodes in the ONTAP environment.

### System components

For the ASA A20, ASA A30, and ASA A50 storage systems, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
Drive	A drive is a device that provides the physical storage needed for data.
Fan	A fan cools the controller and drives.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
NV battery	The non-volatile memory (NV) battery is responsible for providing power to the NVMEM components while data in-flight is being destaged to flash memory after a power loss.

**Power supply** A power supply provides a redundant power source in a controller.

**Real-time clock battery** A real-time clock battery preserves system date and time information if the power is off.

## Boot media

### Boot media replacement workflow - ASA A20, ASA A30, and ASA A50

Get started with replacing the boot media in your ASA A30, ASA A20, or ASA A50 storage system by reviewing the replacement requirements, shutting down the impaired controller, replacing the boot media, restoring the image on the boot media, and verifying the system functionality.

**1**

#### Review the boot media requirements

Review the requirements for boot media replacement.

**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 boot media

Remove the failed boot media from the impaired controller and install the replacement boot media.

**4**

#### Restore the image on the boot media

Restore the ONTAP image from the healthy 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 - ASA A20, ASA A30, and ASA A50

Before replacing the boot media in your ASA A20, ASA A30 or ASA A50 storage system, ensure you meet the necessary requirements and considerations for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0M (wrench) port on the impaired controller is working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

Review the following requirements.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

#### What's next

After you've reviewed the boot media requirements, you [shut down the impaired controller](#).

#### Shut down the controller to replace the boot media - ASA A20, ASA A30, and ASA A50

Shut down the impaired controller in your ASA A20, ASA A30, or ASA A50 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, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

#### About this task

- If you have a SAN system, you must have checked event messages (cluster kernel-service show) for the impaired controller SCSI blade. The cluster kernel-service show command (from priv advanced mode) displays the node name, [quorum status](#) of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

#### Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- Enter *y* when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> The <i>-halt true</i> parameter brings you to the LOADER prompt.

### What's next

After you shut down the impaired controller, you [replace the boot media](#).

### Replace the boot media - ASA A20, ASA A30, and ASA A50

The boot media in your ASA A20, ASA A30, or ASA A50 storage system stores essential firmware and configuration data. The replacement process involves removing the controller module, removing the impaired boot media, installing the replacement boot media, and then reinstalling the controller module.

### About this task

If needed, you can turn on the platform chassis location (blue) LEDs to aid in physically locating the affected platform. Log into the BMC using SSH and enter the `system location-led on` command.

A platform chassis has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

#### Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

## Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

## Steps

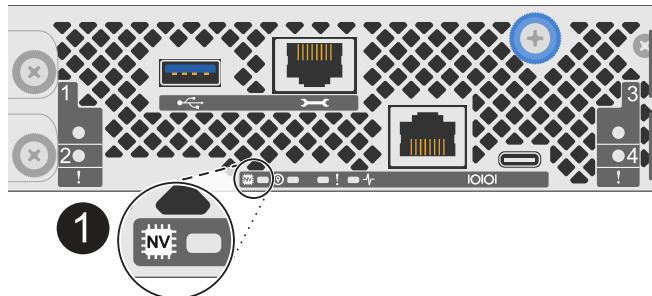
1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.

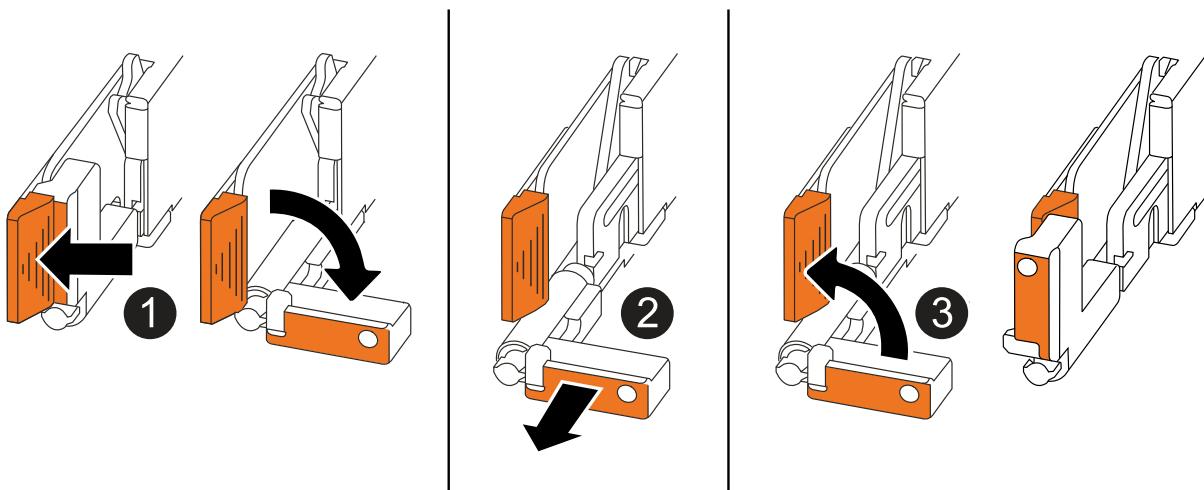
If you are disconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"><li>1. Open the power cord retainer.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>
DC PSU	<ol style="list-style-type: none"><li>1. Unscrew the two thumb screws on the D-SUB DC power cord connector.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	<ul style="list-style-type: none"> <li>Pull the handles towards you to unseat the controller from the midplane.</li> </ul> <p>As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling.</p> <ul style="list-style-type: none"> <li>Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.</li> </ul>
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

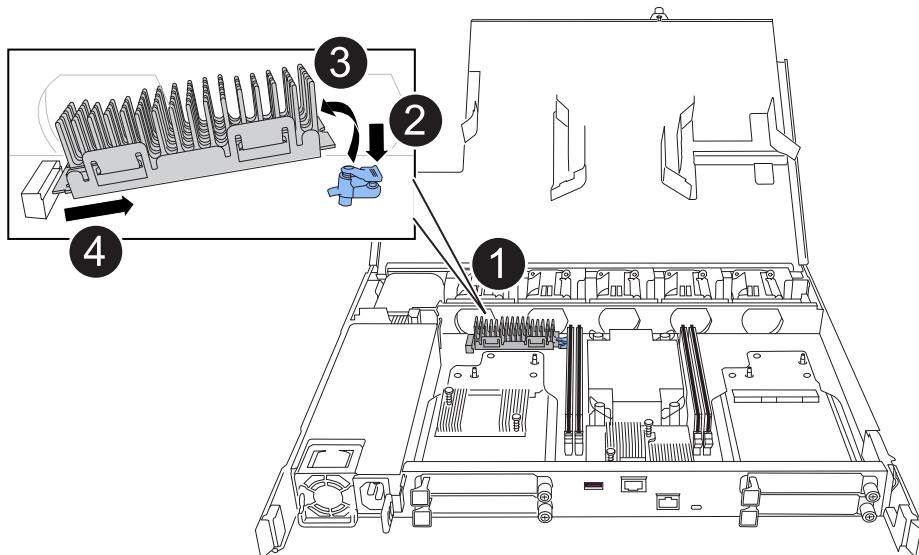
6. Place the controller on an anti-static mat.

7. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

#### Step 2: Replace the boot media

To replace the boot media, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Remove the boot media:



1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

3. Install the replacement boot media:

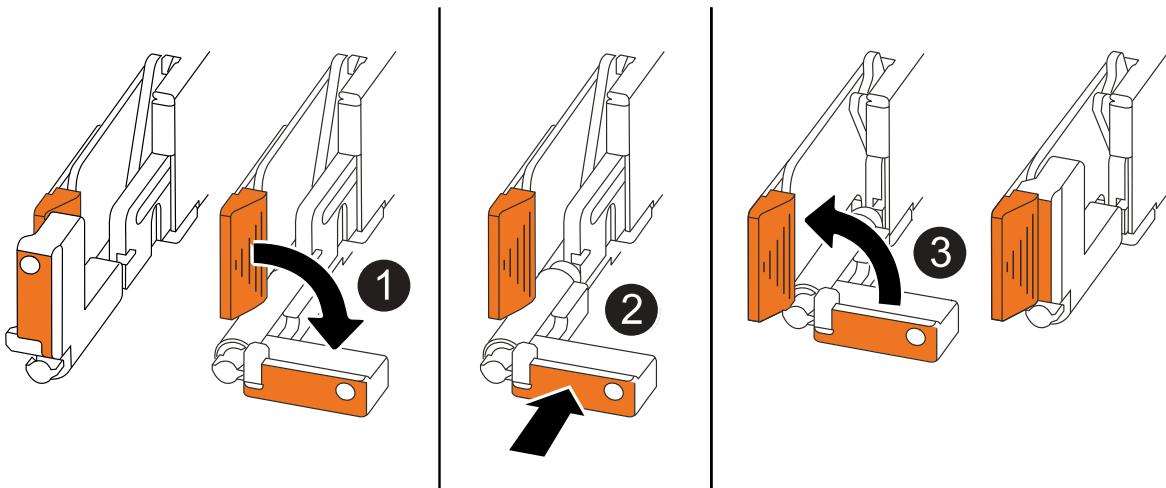
- Remove the boot media from its package.
- Slide the socket end of the boot media into its socket.
- At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

**Step 3: Reinstall the controller**

Reinstall the controller into the chassis and reboot it.

**About this task**

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

## Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.

2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



Do not completely insert the controller in the chassis until instructed to do so later in this procedure.

3. Reconnect the cables to the controller; however, do not plug in the power cord to the power supply (PSU) at this time.



Make sure that the console cable is connected to the controller because you want to catch and log the boot sequence later in the boot media replacement procedure when you fully seat the controller in the chassis and it begins to boot.

4. Fully seat the controller in the chassis:

a. Firmly push on the handles until the controller meets the midplane and is fully seated.

Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.



The controller boots to the LOADER prompt when fully seated in the chassis. It gets its power from the partner controller.

b. Rotate the controller handles up and lock in place with the tabs.

5. Reconnect the power cord to the PSU on the impaired controller.

Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"> <li>1. Plug the power cord into the PSU.</li> <li>2. Secure the power cord with the power cord retainer.</li> </ol>
DC PSU	<ol style="list-style-type: none"> <li>1. Plug the D-SUB DC power cord connector into the PSU.</li> <li>2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.</li> </ol>

### What's next

After physically replacing the impaired boot media, you [restore the ONTAP image from the partner node](#).

### Restore the ONTAP image on the boot media - ASA A20, ASA A30, and ASA A50

After installing the new boot media device in your ASA A20, ASA A30, or ASA A50 storage system, you can start the automated boot media recovery process to restore the configuration from the healthy node.

During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

### Before you begin

- Determine your key manager type:
  - Onboard Key Manager (OKM): Requires cluster-wide passphrase and backup data
  - External Key Manager (EKM): Requires the following files from the partner node:
    - /cfcard/kmip/servers.cfg
    - /cfcard/kmip/certs/client.crt
    - /cfcard/kmip/certs/client.key
    - /cfcard/kmip/certs/CA.pem

### Steps

1. From the LOADER prompt, start the boot media recovery process:

```
boot_recovery -partner
```

The screen displays the following message:

```
Starting boot media recovery (BMR) process. Press Ctrl-C to abort...
```

2. Monitor the boot media install recovery process.

The process completes and displays the Installation complete message.

3. The system checks for encryption and displays one of the following messages:

If you see this message...	Do this...
key manager is not configured. Exiting.	<p>Encryption is not installed on the system.</p> <ol style="list-style-type: none"><li>Wait for the login prompt to display.</li><li>Log into the node and give back the storage:</li></ol> <pre>storage failover giveback -ofnode impaired_node_name</pre> <ol style="list-style-type: none"><li>Go to <a href="#">re-enabling automatic giveback</a> if it was disabled.</li></ol>
key manager is configured.	Encryption is installed. Go to <a href="#">restoring the key manager</a> .



If the system cannot identify the key manager configuration, it displays an error message and prompts you to confirm whether key manager is configured and which type (onboard or external). Answer the prompts to proceed.

4. Restore the key manager using the appropriate procedure for your configuration:

## Onboard Key Manager (OKM)

The system displays the following message and begins running BootMenu Option 10:

```
key manager is configured.  
Entering Bootmenu Option 10...
```

```
This option must be used only in disaster recovery procedures. Are  
you sure? (y or n):
```

- a. Enter **y** at the prompt to confirm you want to start the OKM recovery process.
- b. Enter the passphrase for onboard key management when prompted.
- c. Enter the passphrase again when prompted to confirm.
- d. Enter the backup data for onboard key manager when prompted.

### Show example of passphrase and backup data prompts

```
Enter the passphrase for onboard key management:  
-----BEGIN PASSPHRASE-----  
<passphrase_value>  
-----END PASSPHRASE-----  
Enter the passphrase again to confirm:  
-----BEGIN PASSPHRASE-----  
<passphrase_value>  
-----END PASSPHRASE-----  
Enter the backup data:  
-----BEGIN BACKUP-----  
<passphrase_value>  
-----END BACKUP-----
```

- e. Monitor the recovery process as it restores the appropriate files from the partner node.

When the recovery process is complete, the node reboots. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....  
Setting recovery material for the onboard key manager  
Recovery secrets set successfully  
Trying to delete any existing km_onboard.keydb file.  
  
Successfully recovered keymanager secrets.
```

- f. After the node reboots, verify that the system is back online and operational.

g. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

h. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster:

```
security key-manager onboard sync
```

Go to [re-enabling automatic giveback](#) if it was disabled.

### External Key Manager (EKM)

The system displays the following message and begins running BootMenu Option 11:

```
key manager is configured.  
Entering Bootmenu Option 11...
```

a. Enter the EKM configuration settings when prompted:

i. Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file:

#### Show example of client certificate contents

```
-----BEGIN CERTIFICATE-----  
<certificate_value>  
-----END CERTIFICATE-----
```

ii. Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file:

#### Show example of client key file contents

```
-----BEGIN RSA PRIVATE KEY-----  
<key_value>  
-----END RSA PRIVATE KEY-----
```

iii. Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file:

#### Show example of KMIP server file contents

```
-----BEGIN CERTIFICATE-----  
<KMIP_certificate_CA_value>  
-----END CERTIFICATE-----
```

iv. Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file:

**Show example of server configuration file contents**

```
xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx
xxx.xxx.xxx.xxx:5696.port=5696
xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem
xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4
1xxx.xxx.xxx.xxx:5696.timeout=25
xxx.xxx.xxx.xxx:5696.nbio=1
xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt
xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key
xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:
!RC2:!RC4:!SEED:!eNULL:!aNULL"
xxx.xxx.xxx.xxx:5696.verify=true
xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>
```

v. If prompted, enter the ONTAP Cluster UUID from the partner node. You can check the cluster UUID from the partner node using the `cluster identify show` command.

**Show example of ONTAP Cluster UUID prompt**

```
Notice: bootarg.mgwd.cluster_uuid is not set or is empty.
Do you know the ONTAP Cluster UUID? {y/n} y
Enter the ONTAP Cluster UUID: <cluster_uuid_value>
```

```
System is ready to utilize external key manager(s).
```

vi. If prompted, enter the temporary network interface and settings for the node:

- The IP address for the port
- The netmask for the port
- The IP address of the default gateway

### Show example of temporary network setting prompts

In order to recover key information, a temporary network interface needs to be configured.

```
Select the network port you want to use (for example,  
'e0a')  
e0M
```

```
Enter the IP address for port : xxx.xxx.xxx.xxx  
Enter the netmask for port : xxx.xxx.xxx.xxx  
Enter IP address of default gateway: xxx.xxx.xxx.xxx  
Trying to recover keys from key servers....  
[discover_versions]  
[status=SUCCESS reason= message=]
```

#### b. Verify the key restoration status:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored. The process restores the appropriate files from the partner node and reboots the node. Proceed to the next step.
- If the key is not successfully restored, the system halts and displays error and warning messages. Rerun the recovery process from the **LOADER** prompt: `boot_recovery -partner`

### Show example of key recovery error and warning messages

```
ERROR: kmip_init: halting this system with encrypted
mroot...
WARNING: kmip_init: authentication keys might not be
available.
*****
*          A T T E N T I O N          *
*
*      System cannot connect to key managers.      *
*
*****
ERROR: kmip_init: halting this system with encrypted
mroot...
.
Terminated

Uptime: 11m32s
System halting...

LOADER-B>
```

- c. After the node reboots, verify that the system is back online and operational.
- d. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

Go to [re-enabling automatic giveback](#) if it was disabled.

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

#### What's next

After you've restored the ONTAP image and the node is up and serving data, you [return the failed part to NetApp](#).

#### **Return the failed part to NetApp - ASA A20, ASA A30, and ASA A50**

When a component in your ASA A20, ASA A30, or ASA A50 storage system fails, return the failed part to NetApp. See the [Part Return and Replacements](#) page for further

information.

## Chassis

### Chassis replacement workflow - ASA A20, ASA A30, and ASA A50

Get started with replacing the chassis of your ASA A20, ASA A30, or ASA A50 storage system by reviewing the replacement requirements, shutting down the controllers, replacing the chassis, and verifying system operations.

1

#### Review the chassis replace requirements

Review the requirements to replace the chassis.

2

#### Shut down the controllers

Shut down the controllers so you can perform maintenance on the chassis.

3

#### Replace the chassis

Replace the chassis by moving the drives and any drive blanks, controllers (with the power supplies), and bezel from the impaired chassis to the new chassis, and swapping out the impaired chassis with the new chassis of the same model as the impaired chassis.

4

#### Complete chassis replacement

Verify the HA state of the chassis and return the failed part to NetApp.

### Requirements to replace the chassis - ASA A20, ASA A30, and ASA A50

Before replacing the chassis of your ASA A20, ASA A30, or ASA A50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement chassis, and the necessary tools.

Review the following requirements and considerations.

#### Requirements

- The replacement chassis must be the same model as the impaired chassis. This procedure is for a like-for-like replacement, not for an upgrade.
- All other components in the storage system must be functioning properly; if not, contact [NetApp Support](#) before continuing with this procedure.

#### Considerations

- The chassis replacement procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

- You can use the chassis replacement procedure with all versions of ONTAP supported by your storage system.
- The chassis replacement procedure is written with the assumption that you are moving the bezel, drives, any drive blanks, and controllers to the new chassis.

## What's next?

After you've reviewed the requirements to replace the chassis, you need to [shut down the controllers](#).

### Shut down the controllers to replace the chassis - ASA A20, ASA A30, and ASA A50

Shut down the controllers in your ASA A20, ASA A30, or ASA A50 storage system to prevent data loss and ensure system stability when replacing the chassis.

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see [Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base](#).

#### Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessibility for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional [system health checks](#).
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any [Active IQ Wellness Alerts and Risks](#).
 Make note of any faults presently on the system, such as LEDs on the system components.

#### Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

```
system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
```

5. Identify the SP/BMC address of all cluster nodes:

```
system service-processor show -node * -fields address
```

6. Exit the cluster shell:

```
exit
```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous

step to monitor progress.

If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

#### 8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown  
true -ignore-quorum-warnings true -inhibit-takeover true
```



For clusters using SnapMirror synchronous operating in StrictSync mode: system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict -sync-warnings true

#### 9. Enter **y** for each controller in the cluster when you see:

Warning: Are you sure you want to halt node <node\_name>? {y|n}:

#### 10. Wait for each controller to halt and display the LOADER prompt.

#### What's next?

After you've shut down the controllers, you need to [replace the chassis](#).

### Replace the chassis - ASA A20, ASA A30, and ASA A50

Replace the chassis of your ASA A20, ASA A30, or ASA A50 storage system when a hardware failure requires it. The replacement process involves removing the controllers, removing the drives, installing the replacement chassis, and reinstalling the chassis components.

#### About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

#### Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

#### Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

#### Steps

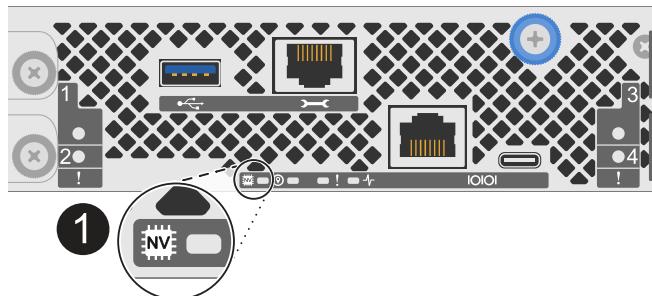
1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1

NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.

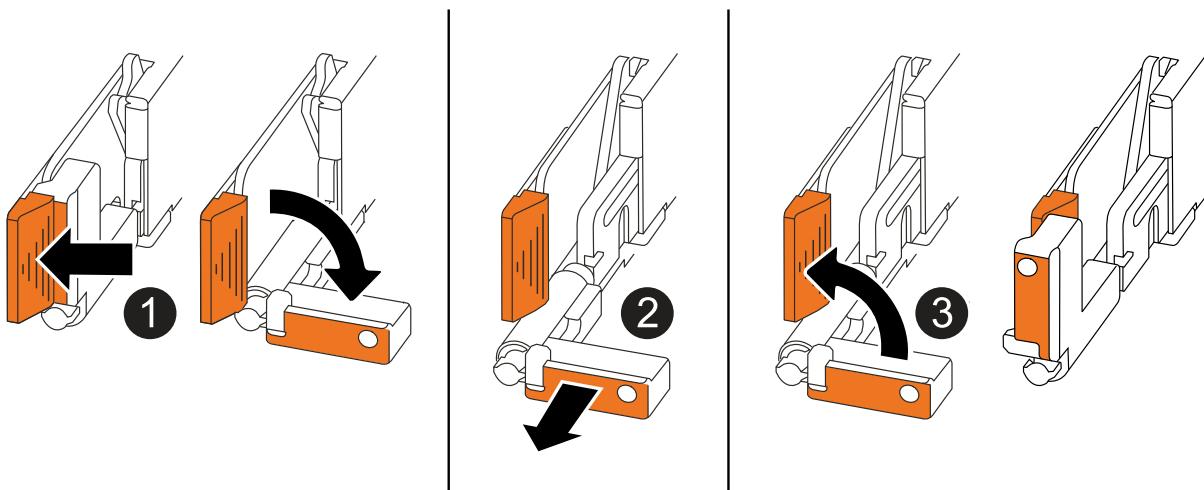
If you are disconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"><li>1. Open the power cord retainer.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>
DC PSU	<ol style="list-style-type: none"><li>1. Unscrew the two thumb screws on the D-SUB DC power cord connector.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	<ul style="list-style-type: none"> <li>Pull the handles towards you to unseat the controller from the midplane.</li> </ul> <p>As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling.</p> <ul style="list-style-type: none"> <li>Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.</li> </ul>
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Repeat these steps for the other controller in the chassis.

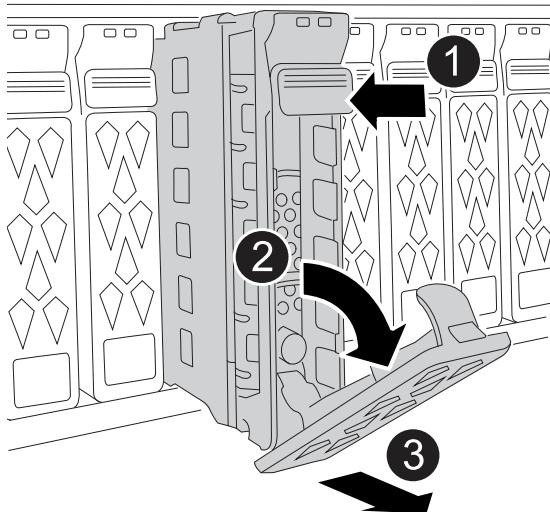
#### Step 2: Remove the drives from the impaired chassis

You need to remove all of the drives and any drive blanks from the impaired chassis so that later in the procedure you can install them in the replacement chassis.

1. Gently remove the bezel from the front of the storage system.
2. Remove the drives and any drive blanks:



Keep track of what drive bay each drive and drive blank was removed from because they must be installed in the same drive bays in the replacement chassis.



1	Press the release button on the drive face to open the cam handle.
2	Rotate the cam handle downward to disengage the drive from the midplane.
3	<p>Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand.</p> <p>When removing a drive, always use two hands to support its weight.</p> <p><span style="border: 1px solid #ccc; border-radius: 50%; padding: 2px 5px; margin-right: 10px;">i</span> Because drives are fragile, minimize handling to avoid damaging them.</p>

3. Set the drives aside on a static-free cart or table.

#### Step 2: Replace the chassis from within the equipment rack or system cabinet

You remove the impaired chassis from the equipment rack or system cabinet, install the replacement chassis, install the drives, any drive blanks, and then install the bezel.

1. Remove the screws from the impaired chassis mount points.

Set the screws aside to use later in this procedure.



If the storage system shipped in a NetApp system cabinet, you must remove additional screws at the rear of the chassis before the chassis can be removed.

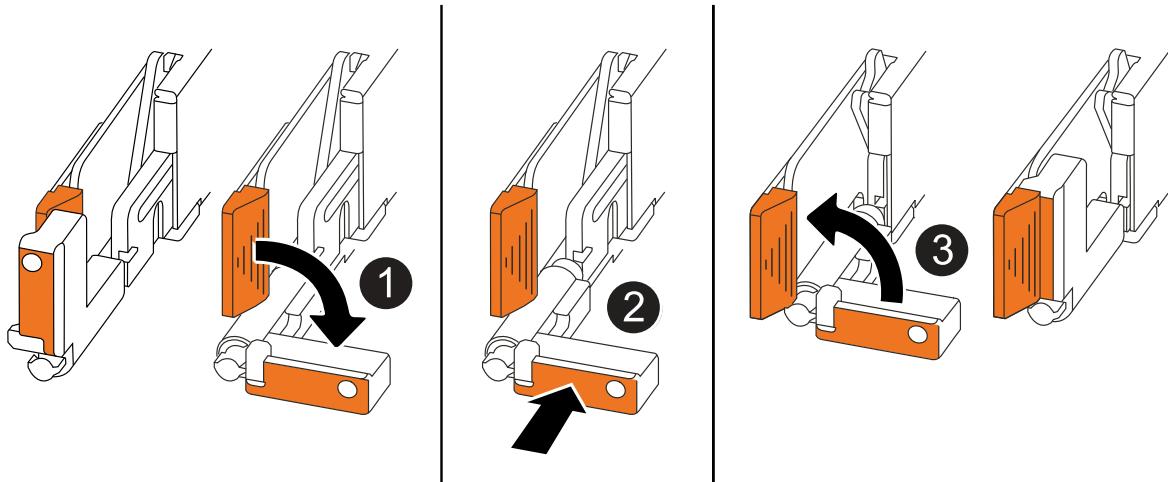
2. Using two people or a power lift, remove the impaired chassis from the equipment rack or system cabinet by sliding it off the rails, and then set it aside.
3. Using two people, install the replacement chassis into the equipment rack or system cabinet by sliding it onto the rails.
4. Secure the front of the replacement chassis to the equipment rack or system cabinet using the screws you removed from the impaired chassis.

#### Step 4: Install the controllers and drives

Install the controllers and drives into the replacement chassis and reboot the controllers.

##### About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when installing a controller, and can be used as a reference for the rest of the controller installation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis and push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

1. Insert one of the controllers into the chassis:

- Align the back of the controller with the opening in the chassis.
- Firmly push on the handles until the controller meets the midplane and is fully seated in the chassis.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

- Rotate the controller handles up and lock in place with the tabs.

2. Recable the controller, as needed, except for the power cords.

3. Repeat these steps to install the second controller into the chassis.

4. Install the drives and any drive blanks you removed from the impaired chassis into the replacement chassis:



The drives and drive blanks must be installed in the same drive bays in the replacement chassis.

- With the cam handle in the open position, use both hands to insert the drive.

- b. Gently push until the drive stops.
- c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

- d. Repeat the process for the remaining drives.

5. Install the bezel.

6. Reconnect the power cords to the power supplies (PSU) in the controllers.

Once power is restored to a PSU, the status LED should be green.



The controllers begin to boot as soon as the power is restored.

If you are reconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"> <li>1. Plug the power cord into the PSU.</li> <li>2. Secure the power cord with the power cord retainer.</li> </ol>
DC PSU	<ol style="list-style-type: none"> <li>1. Plug the D-SUB DC power cord connector into the PSU.</li> <li>2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.</li> </ol>

7. If controllers boot to the LOADER prompt, reboot the controllers:

`boot_ontap`

8. Turn AutoSupport back on:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

### Complete chassis replacement - ASA A20, ASA A30, and ASA A50

Verify the HA state of the chassis and then return the failed part to NetApp to complete the final step in the ASA A20, ASA A30, and ASA A50 chassis replacement procedure.

#### Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your storage system configuration.

1. In Maintenance mode, from either controller, display the HA state of the local controller and chassis:

```
ha-config show
```

The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your storage system configuration:

- a. Set the HA state for the chassis:

```
ha-config modify chassis HA-state
```

The value for HA-state should be *ha*.

The value for HA-state can be one of the following:

- \* **ha**
- \* *mcc* (not supported in ASA)

b. Confirm that the setting has changed:

```
ha-config show
```

3. If you have not already done so, recable the rest of your storage system.

## Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

## Controller

### Controller replacement workflow - ASA A20, ASA A30, and ASA A50

Get started with replacing the controller in your ASA A20, ASA A30, or ASA A50 storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.

1

#### [Review the controller replacement requirements](#)

Review the requirements to replace the controller.

2

#### [Shut down the impaired controller](#)

Shut down or take over the impaired controller so that the healthy controller continues to serve data from the impaired controller storage.

3

#### [Replace the controller](#)

Replacing the controller includes removing the impaired controller, moving FRU components to the replacement controller, installing the replacement controller in the chassis, setting the time and date, and then recabling.

4

#### [Restore and verify the system configuration](#)

Verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

5

#### [Give back the controller](#)

Transfer the ownership of storage resources back to the replacement controller.

## 6

### Complete controller replacement

Verify the LIFs, check cluster health, and return the failed part to NetApp.

#### Requirements to replace the controller - ASA A20, ASA A30, and ASA A50

Before replacing the controller in your ASA A20, ASA A30, or ASA A50 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.

You must review the requirements and considerations for the controller replacement procedure.

#### Requirements

- All shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- You must replace a controller with a controller of the same model type. You cannot upgrade your system by just replacing the controller.
- You cannot change any drives or shelves as part of this procedure.
- You must always capture the controller's console output to a text log file.

The console output provides you with a record of the procedure you can use to troubleshoot issues you might encounter during the replacement process.

#### Considerations

- It is important that you apply the commands in this procedure to the correct controller:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.

#### What's next?

After you've reviewed the requirements to replace the impaired controller, you need to [shut down the impaired controller](#).

#### Shut down the impaired controller - ASA A20, ASA A30, and ASA A50

Shut down the impaired controller in your ASA A20, ASA A30, or ASA A50 storage system to prevent data loss and ensure system stability when replacing the controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

## About this task

- If you have a SAN system, you must have checked event messages (cluster kernel-service show) for the impaired controller SCSI blade. The cluster kernel-service show command (from priv advanced mode) displays the node name, **quorum status** of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

## Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter *y* when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller:  <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> The <i>-halt true</i> parameter brings you to the LOADER prompt.

## What's next?

After you've shut down the impaired controller, you need to [replace the controller](#).

## Replace the controller - ASA A20, ASA A30, and ASA A50

Replace the controller in your ASA A20, ASA A30, or ASA A50 storage system when a

hardware failure requires it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

#### About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

#### Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

#### Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

#### Steps

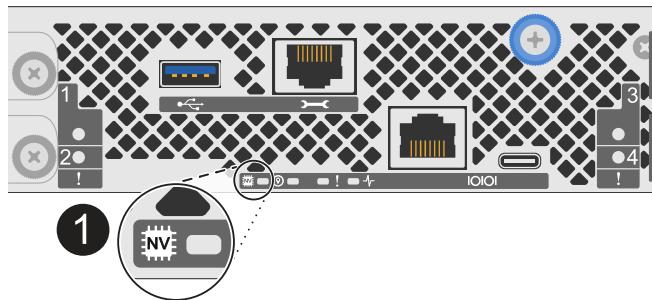
1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.

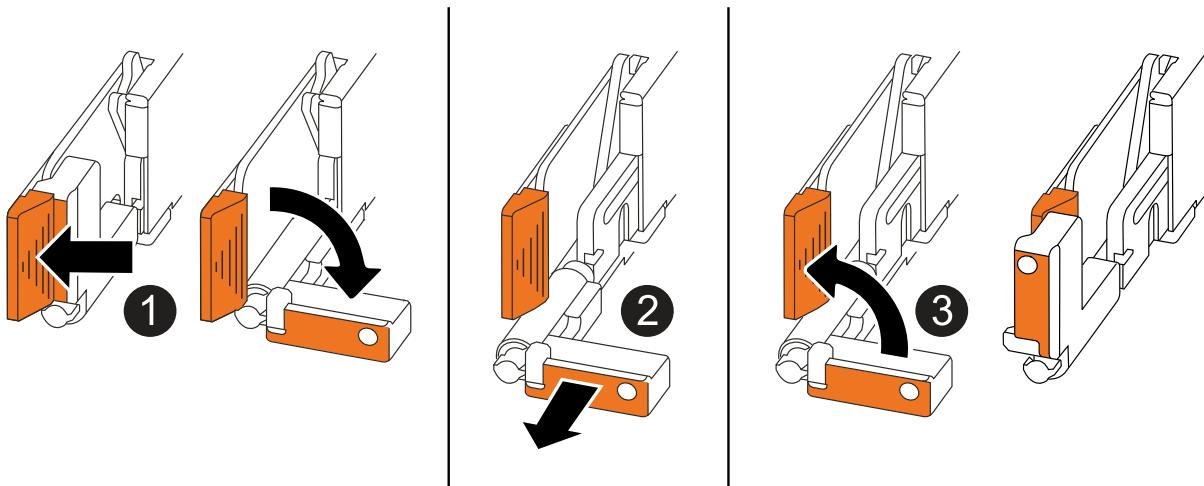
If you are disconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"><li>1. Open the power cord retainer.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>
DC PSU	<ol style="list-style-type: none"><li>1. Unscrew the two thumb screws on the D-SUB DC power cord connector.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	<ul style="list-style-type: none"><li>• Pull the handles towards you to unseat the controller from the midplane.</li></ul> <p>As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling.</p> <ul style="list-style-type: none"><li>• Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.</li></ul>
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

## **Step 2: Move the power supply**

Move the power supply (PSU) to the replacement controller.

1. Move the PSU from the impaired controller:

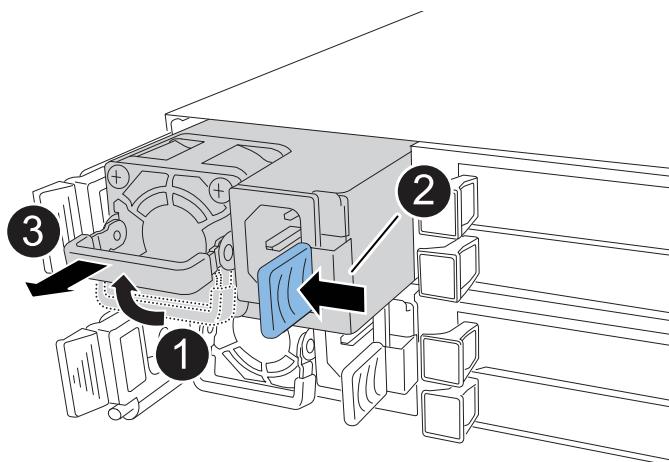
Make sure the left side controller handle is in the upright position to allow you access to the PSU.

## Option 1: Move an AC PSU

To move an AC PSU, complete the following steps.

### Steps

1. Remove the AC PSU from the impaired controller:



1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	Pull the PSU out of the controller while using your other hand to support its weight.   The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.

2. Insert the PSU into the replacement controller:

- a. Using both hands, support and align the edges of the PSU with the opening in the controller.
- b. Gently push the PSU into the controller until the locking tab clicks into place.

A PSU will only properly engage with the internal connector and lock in place one way.



To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

- c. Rotate the handle down, so it is out of the way of normal operations.

## Option 2: Move a DC PSU

To move a DC PSU, complete the following steps.

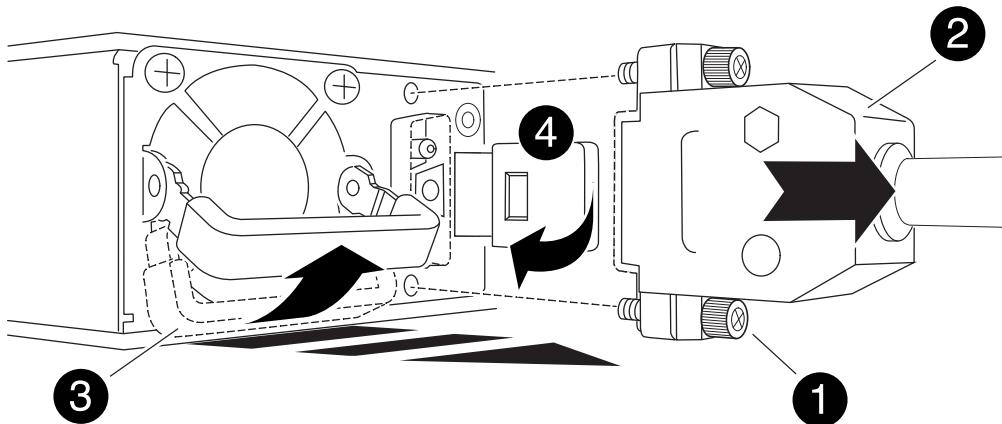
### Steps

1. Remove the DC PSU from the impaired controller:

- a. Rotate the handle up, to its horizontal position, and then grasp it.
- b. With your thumb, press the terracotta tab to release the locking mechanism.
- c. Pull the PSU out of the controller while using your other hand to support its weight.



The PSU is short. Always use two hands to support it when removing it from the controller so that it does not swing free from the controller and injure you.



1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

2. Insert the PSU into the replacement controller:

- a. Using both hands, support and align the edges of the PSU with the opening in the controller.
- b. Gently slide the PSU into the controller until the locking tab clicks into place.

A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.



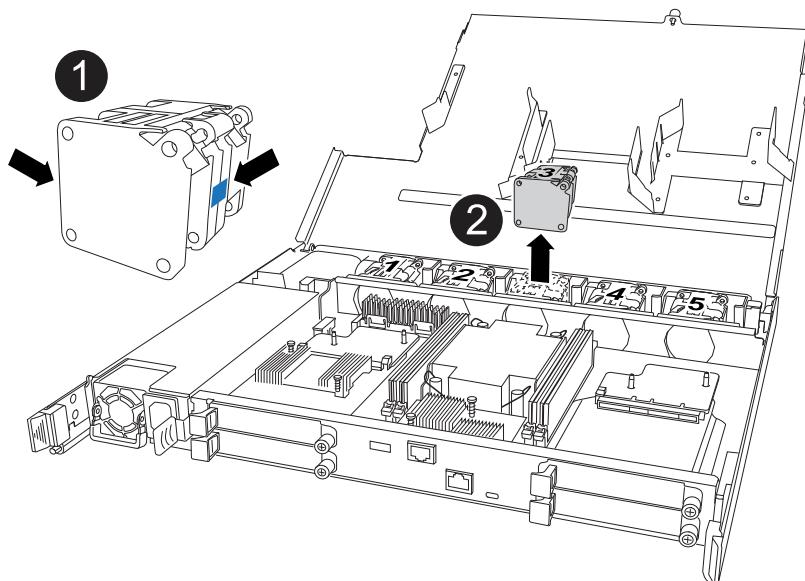
To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

- c. Rotate the handle down, so it is out of the way of normal operations.

### Step 3: Move the fans

Move the fans to the replacement controller.

1. Remove one of the fans from the impaired controller:



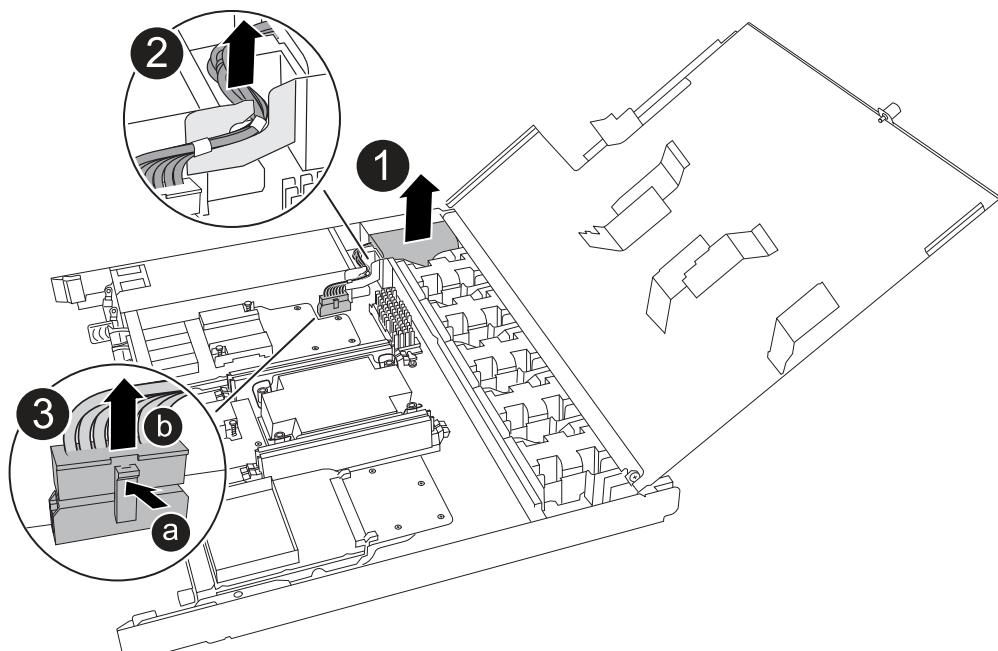
1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

2. Insert the fan into the replacement controller by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.
3. Repeat these steps for the remaining fans.

#### Step 4: Move the NV battery

Move the NV battery to the replacement controller.

1. Remove the NV battery from the impaired controller:



1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	<ol style="list-style-type: none"> <li>1. Push in and hold the tab on the connector.</li> <li>2. Pull the connector up and out of the socket.</li> </ol> <p>As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.</p>

2. Install the NV battery into the replacement controller:

- a. Plug the wiring connector into its socket.
- b. Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
- c. Place the NV battery into the compartment.

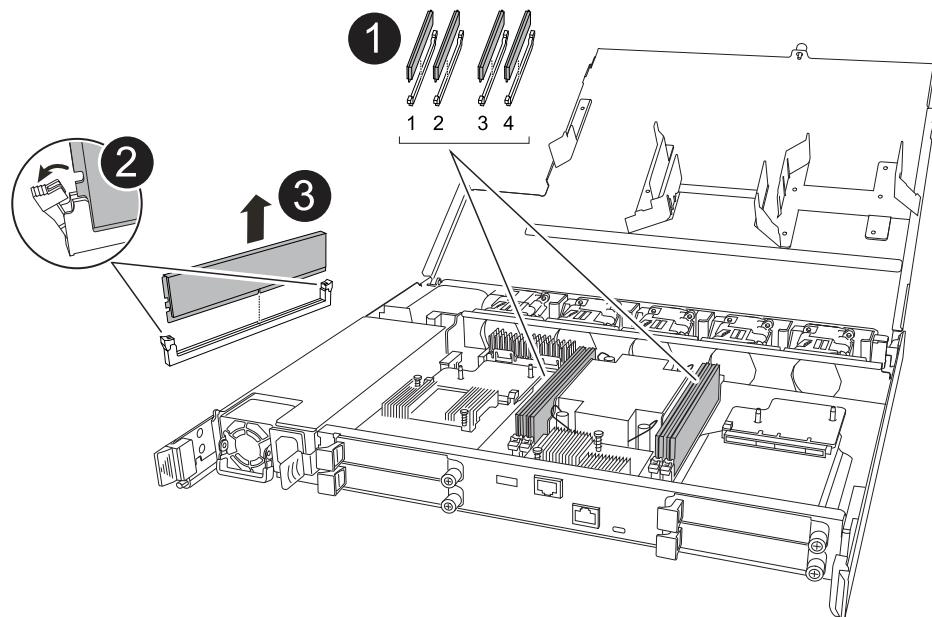
The NV battery should sit flush in its compartment.

**Step 5: Move system DIMMs**

Move the DIMMs to the replacement controller.

If you have DIMM blanks, you do not need to move them, the replacement controller should come with them installed.

1. Remove one of the DIMMs from the impaired controller:



1	<p>DIMM slot numbering and positions.</p> <p> Depending on your storage system model, you will have two or four DIMMs.</p>
2	<ul style="list-style-type: none"> <li>• Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller in the proper orientation.</li> <li>• Eject the DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot.</li> </ul> <p> Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.</p>
3	<p>Lift the DIMM up and out of the slot.</p> <p>The ejector tabs remain in the open position.</p>

2. Install the DIMM in the replacement controller:

- Make sure that the DIMM ejector tabs on the connector are in the open position.
- Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

When inserted correctly, the DIMM goes in easily but fits tightly in the slot. If not, reinsert the DIMM.

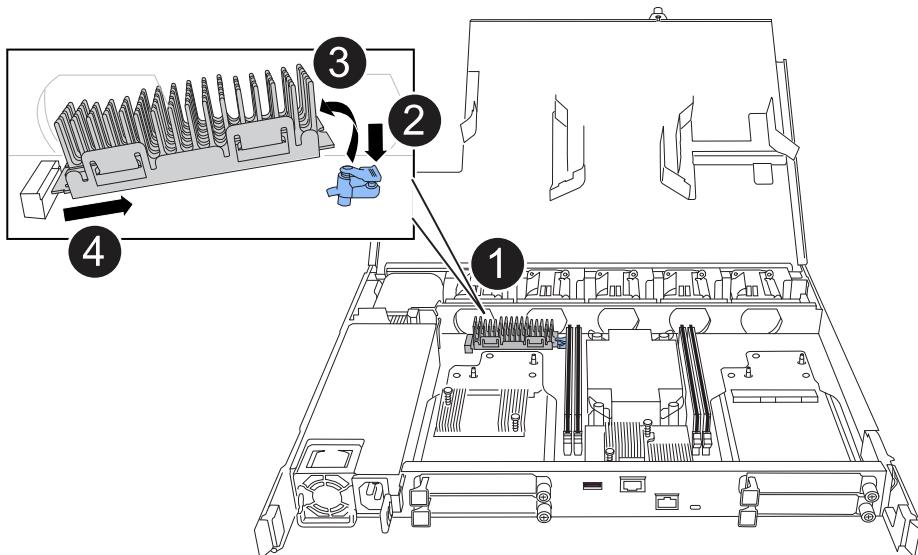
- Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
- Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.

3. Repeat these steps for the remaining DIMMs.

**Step 6: Move the boot media**

Move the boot media to the replacement controller.

1. Remove the boot media from the impaired controller:



1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

2. Install the boot media into the replacement controller:

- Slide the socket end of the boot media into its socket.
- At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

**Step 7: Move the I/O modules**

Move the I/O modules and any I/O blanking modules to the replacement controller.

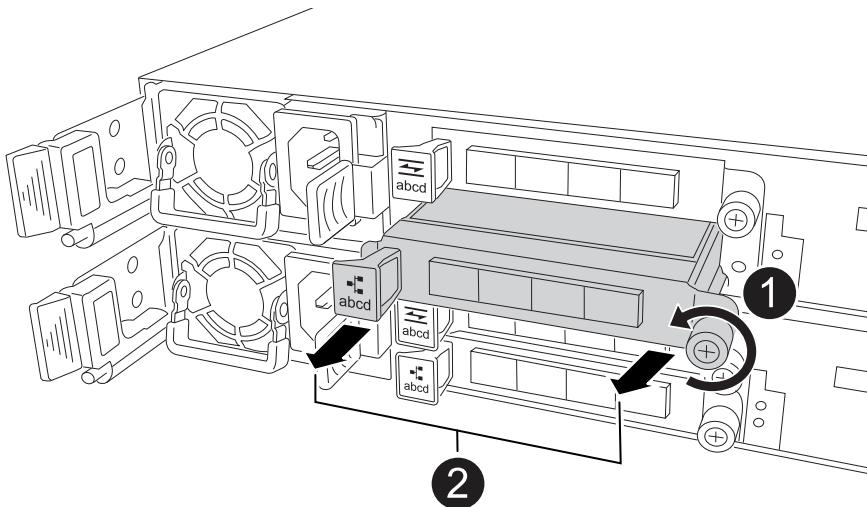
1. Unplug cabling from one of the I/O modules.

Make sure to label the cables so that you know where they came from.

2. Remove the I/O module from the impaired controller:

Make sure that you keep track of which slot the I/O module was in.

If you are removing the I/O module in slot 4, make sure the right side controller handle is in the upright position to allow you access to the I/O module.



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

3. Install the I/O module into the replacement controller:

a. Align the I/O module with the edges of the slot.

b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O Module.

c. Turn the thumbscrew clockwise to tighten.

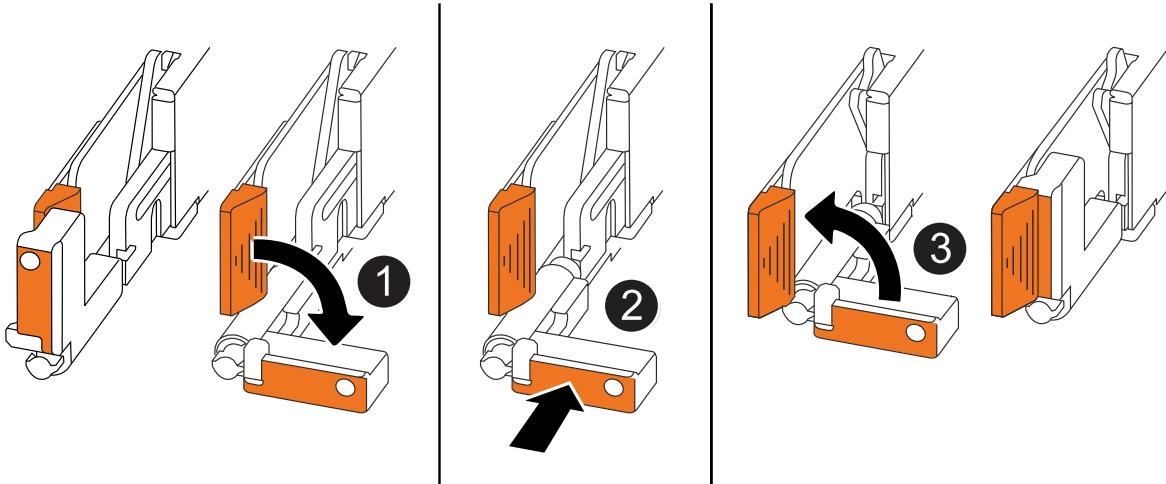
4. Repeat these steps to move the remaining I/O modules and any I/O blanking modules to the replacement controller.

**Step 8: Install the controller**

Reinstall the controller into the chassis and reboot it.

**About this task**

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

## Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.

2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.



Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:

a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

b. Rotate the controller handles up and lock in place with the tabs.



The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Take the controller to the LOADER prompt by pressing CTRL-C to abort AUTOBOOT.

6. Set the time and date on the controller:

Make sure you are at the controller's LOADER prompt.

a. Display the date and time on the controller:

```
show date
```



Time and date default is in GMT. You have the option to display in local time and in 24hr mode.

b. Set the current time in GMT:

```
set time hh:mm:ss
```

You can get the current GMT from the healthy node:

```
date -u
```

c. Set the current date in GMT:

```
set date mm/dd/yyyy
```

You can get the current GMT from the healthy node:

```
date -u
```

7. Recable the controller as needed.

8. Reconnect the power cord to the power supply (PSU).

Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"><li>1. Plug the power cord into the PSU.</li><li>2. Secure the power cord with the power cord retainer.</li></ol>
DC PSU	<ol style="list-style-type: none"><li>1. Plug the D-SUB DC power cord connector into the PSU.</li><li>2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.</li></ol>

#### What's next?

After you've replaced the impaired controller, you need to [restore the system configuration](#).

#### Restore and verify the system configuration - ASA A20, ASA A30, and ASA A50

Verify that the controller's HA configuration is active and functioning correctly in your ASA A20, ASA A30, or ASA A50 storage system, and confirm that the system's adapters list all the paths to the disks.

## Step 1: Verify HA config settings

You must verify the HA state of the controller and, if necessary, update the state to match your storage system configuration.

1. Boot to maintenance mode:

```
boot_ontap maint
```

- a. Enter `y` when you see *Continue with boot?*.

If you see the *System ID mismatch* warning message, enter `y`.

2. Enter `sysconfig -v` and capture the display contents.



If you see *PERSONALITY MISMATCH* contact customer support.

3. From the `sysconfig -v` output, compare the adapter card information with the cards and locations in the replacement controller.

4. Verify that all components display the same HA state:

```
ha-config show
```

The HA state should be the same for all components.

5. If the displayed system state of the controller does not match your storage system configuration, set the HA state for the controller:

```
ha-config modify controller ha
```

The value for the HA state can be one of the following:

- ha
- mcc (not supported)
- mccip (not supported in ASA systems)
- non-ha (not supported)

6. Confirm that the setting has changed:

```
ha-config show
```

## Step 2: Verify disk list

1. Verify that the adapter lists the paths to all disks:

```
storage show disk -p
```

If you see any issues, check cabling and reseat cables.

2. Exit Maintenance mode:

halt

#### What's next?

After you've restored and verified your system configuration, you need to [give back the controller](#).

#### Give back the controller - ASA A20, ASA A30, and ASA A50

Return control of storage resources to the replacement controller so your ASA A20, ASA A30, or ASA A50 storage system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption, Onboard Key Manager (OKM) encryption, or External Key Manager (EKM) encryption.

## No encryption

Return the impaired controller to normal operation by giving back its storage.

### Steps

1. From the LOADER prompt, enter `boot_ontap`.
2. Press `<enter>` when console messages stop.
  - If you see the *login* prompt, go to the next step at the end of this section.
  - If you see *Waiting for giveback*, press the `<enter>` key, log into the partner node, and then go to the next step at the end of this section.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

## Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

### Steps

1. From the LOADER prompt, enter `boot_ontap maint`.
2. Boot to the ONTAP menu from the LOADER prompt `boot_ontap` menu and select option 10.
3. Enter the OKM passphrase.



You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option 1 for normal boot.
6. Press `<enter>` when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as admin.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`
9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
  - a. Move the console cable back to the replacement controller.
  - b. Synchronize missing keys: `security key-manager onboard sync`



If you encounter errors, contact [NetApp Support](#).



You are prompted for the cluster-wide passphrase of OKM for the cluster.

c. Verify status of the keys: `security key-manager key query -restored false`

The output should show no results when properly synchronized.

If the output shows results (the key IDs of keys that are not present in the system's internal key table), contact [NetApp Support](#).

11. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
12. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
13. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

### External key manager (EKM)

Reset encryption and return the controller to normal operation.

#### Steps

1. If the root volume is encrypted with External Key Manager and the console cable is connected to the replacement node, enter `boot_ontap` menu and select option 11.
2. If these questions appear, answer y or n as appropriate:

Do you have a copy of the `/cfcard/kmip/certs/client.crt` file? {y/n}

Do you have a copy of the `/cfcard/kmip/certs/client.key` file? {y/n}

Do you have a copy of the `/cfcard/kmip/certs/CA.pem` file? {y/n}

Do you have a copy of the `/cfcard/kmip/servers.cfg` file? {y/n}

Do you know the KMIP server address? {y/n}

Do you know the KMIP port? {y/n}



Contact [NetApp Support](#) if you have issues.

3. Supply the information for:
  - The client certificate (`client.crt`) file contents
  - The client key (`client.key`) file contents
  - The KMIP server CA(s) (`CA.pem`) file contents
  - The IP address for the KMIP server
  - The port for the KMIP server
4. Once the system processes, you see the Boot Menu. Select '1' for normal boot.
5. Check the takeover status: `storage failover show`
6. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
7. If automatic giveback was disabled, reenable it: `storage failover modify -node local`

```
-auto-giveback true
```

8. If AutoSupport is enabled, restore/unsuppress automatic case creation: system node autosupport invoke -node \* -type all -message MAINT=END

## What's next?

After you've transferred the ownership of storage resources to the replacement controller, you need to [complete the controller replacement](#) procedure.

## Complete controller replacement - ASA A20, ASA A30, and ASA A50

To complete the controller replacement for your ASA A20, ASA A30, or ASA A50 storage system, first restore the NetApp Storage Encryption configuration (if necessary) and install the required licenses on the new controller. Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, register the new controller's serial number and then return the failed part to NetApp.

### Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

#### Before you begin

If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in [Post Motherboard Replacement Process to update Licensing on ONTAP platforms](#). If you are unsure of the initial ONTAP release for your system, see [NetApp Hardware Universe](#) for more information.

#### About this task

- Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

- The license keys must be in the 28-character format.
- You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

#### Steps

1. If you need new license keys, obtain replacement license keys on the [NetApp Support Site](#) in the My Support section under Software licenses.



The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: system license add -license-code license-key... .

3. Remove the old licenses, if desired:

- a. Check for unused licenses: `license clean-up -unused -simulate`
- b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

**Step 2: Verify LIFs, register the serial number, and check cluster health**

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

**Steps**

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.

- If AutoSupport is enabled, send an AutoSupport message to register the serial number.
- If AutoSupport is not enabled, call [NetApp Support](#) to register the serial number.

3. Check the health of your cluster. See the [How to perform a cluster health check with a script in ONTAP](#) KB article for more information.

4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

**Step 3: Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

## Replace a DIMM - ASA A20, ASA A30, and ASA A50

Replace a DIMM in your ASA A20, ASA A30, or ASA A50 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.

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

**Before you begin**

- Ensure all other components in the storage system are working correctly; if not, contact [NetApp Support](#) before continuing."
- You must replace the failed FRU component with a replacement FRU component you received from your provider.

## About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

## Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

### About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, `quorum status` of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

### Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <code>y</code> when prompted.

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

## Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

### Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

### Steps

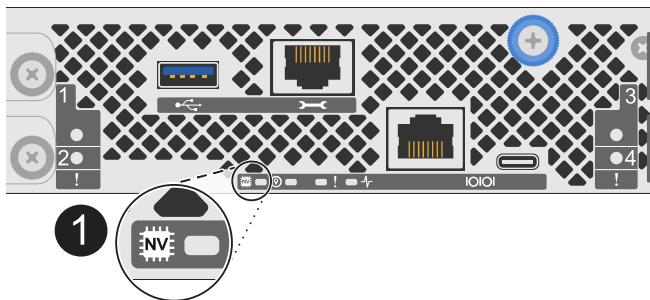
1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1

NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.

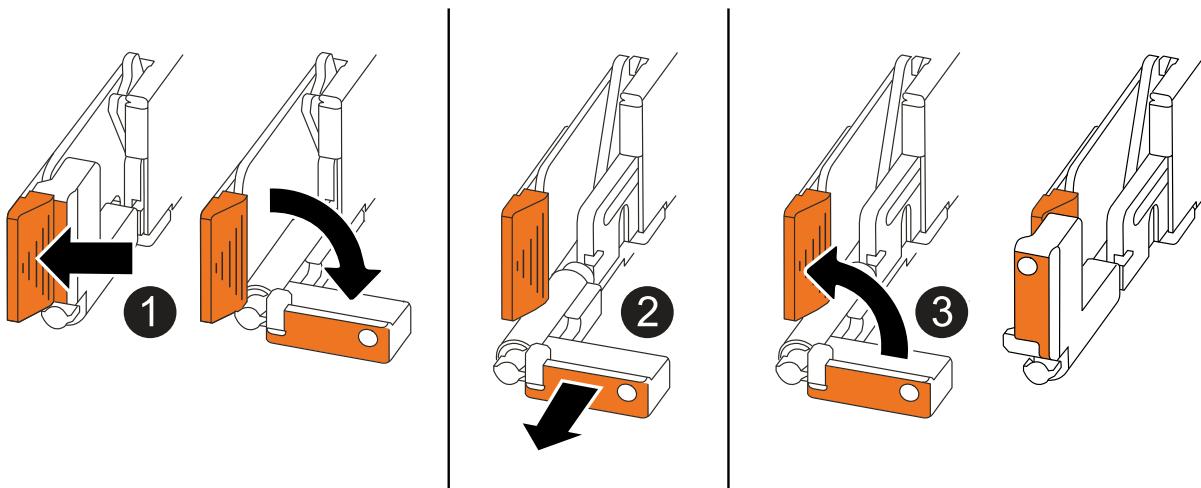
If you are disconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"> <li>1. Open the power cord retainer.</li> <li>2. Unplug the power cord from the PSU and set it aside.</li> </ol>
DC PSU	<ol style="list-style-type: none"> <li>1. Unscrew the two thumb screws on the D-SUB DC power cord connector.</li> <li>2. Unplug the power cord from the PSU and set it aside.</li> </ol>

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	<ul style="list-style-type: none"> <li>• Pull the handles towards you to unseat the controller from the midplane.</li> </ul> <p>As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling.</p> <ul style="list-style-type: none"> <li>• Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.</li> </ul>
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

### Step 3: Replace a DIMM

To replace a DIMM, locate the faulty DIMM inside the controller and follow the specific sequence of steps.

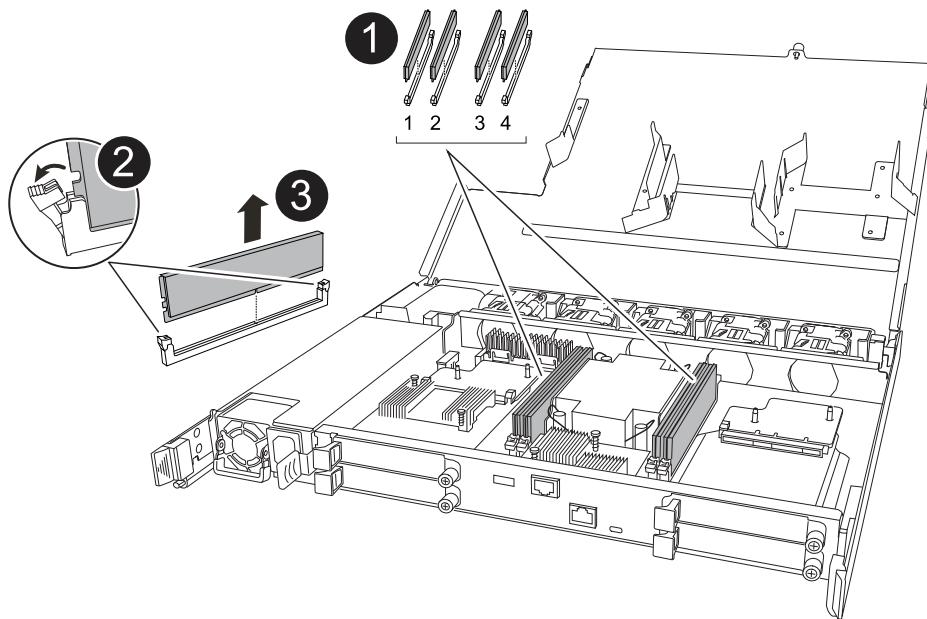
#### Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller and identify the faulty DIMM.



Consult either the [Netapp Hardware Universe](#) or the FRU map on the cover of the controller for exact DIMM locations.

3. Remove the faulty DIMM:



<b>1</b>	<p>DIMM slot numbering and positions.</p> <p> Depending on your storage system model you will have two or four DIMMs.</p>
<b>2</b>	<ul style="list-style-type: none"><li>• Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM using the same orientation.</li><li>• Eject the faulty DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot.</li></ul> <p> Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.</p>
<b>3</b>	<p>Lift the DIMM up and out of the slot.</p> <p>The ejector tabs remain in the open position.</p>

#### 4. Install the replacement DIMM:

- Remove the replacement DIMM from its antistatic shipping bag.
- Make sure that the DIMM ejector tabs on the connector are in the open position.
- Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

When inserted correctly, the DIMM goes in easily but fits tightly in the slot. Reinsert the DIMM if you feel it is not inserted correctly.

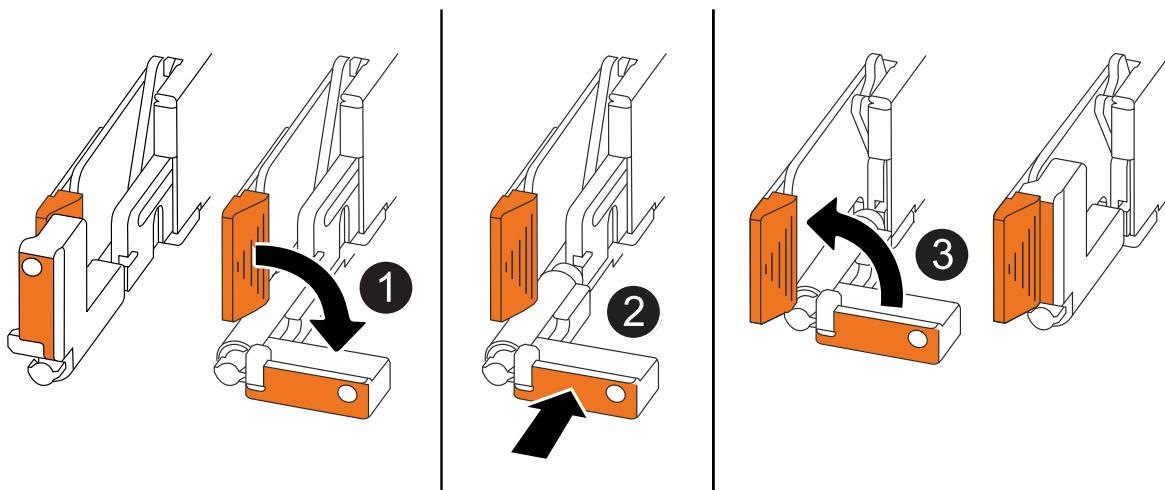
- Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
- Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.

### Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

#### About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

#### Steps

- Close the controller cover and turn the thumbscrew clockwise until tightened.
- Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.



Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:

- a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

- b. Rotate the controller handles up and lock in place with the tabs.



The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.

6. Reconnect the power cord to the power supply (PSU).

Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"><li>1. Plug the power cord into the PSU.</li><li>2. Secure the power cord with the power cord retainer.</li></ol>
DC PSU	<ol style="list-style-type: none"><li>1. Plug the D-SUB DC power cord connector into the PSU.</li><li>2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.</li></ol>

7. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

8. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

## Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

## Replace a Drive - ASA A20, ASA A30, and ASA A50

Replace a drive in your ASA A20, ASA A30, or ASA A50 storage system when a drive fails or requires an upgrade. The replacement process involves identifying the faulty drive, safely removing it, and installing a new drive to ensure continued data access and system performance.

You can replace a failed drive nondisruptively while I/O is in progress.

### Before you begin

- The drive that you are installing must be supported by your storage system.

#### [NetApp Hardware Universe](#)

- If self-encrypting drive (SED) authentication is enabled, you must use the SED replacement instructions in the ONTAP documentation.

Instructions in the ONTAP documentation describe additional steps you must perform before and after replacing an SED.

#### [NetApp encryption overview with the CLI](#)

- All other components in the storage system must be functioning properly, if not, you must contact [NetApp Support](#) before continuing with this procedure.
- Verify that the drive you are removing is failed.

You can verify that the drive is failed by running the `storage disk show -broken` command. The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.



Depending on the drive type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

### About this task

- When replacing a failed drive, you must wait 70 seconds between the removal of the drive and the insertion of the replacement drive to allow the storage system to recognize that a drive was removed.
- The best practice is to have the current version of the Disk Qualification Package (DQP) installed before hot-swapping a drive.

Having the current version of the DQP installed allows your system to recognize and use newly qualified drives. This avoids system event messages about having noncurrent drive information and prevention of drive partitioning because drives are not recognized. The DQP also notifies you of noncurrent drive firmware.

#### [NetApp Downloads: Disk Qualification Package](#)

- The best practice is to have current versions of NVMe shelf module (NSM) firmware and drive firmware on your system before replacing FRU components.

[NetApp Downloads: Disk Shelf Firmware](#)

[NetApp Downloads: Disk Drive Firmware](#)



Do not revert firmware to a version that does not support your shelf and its components.

- Drive firmware is automatically updated (nondisruptively) on new drives that have non-current firmware versions.



Drive firmware checks occur every two minutes.

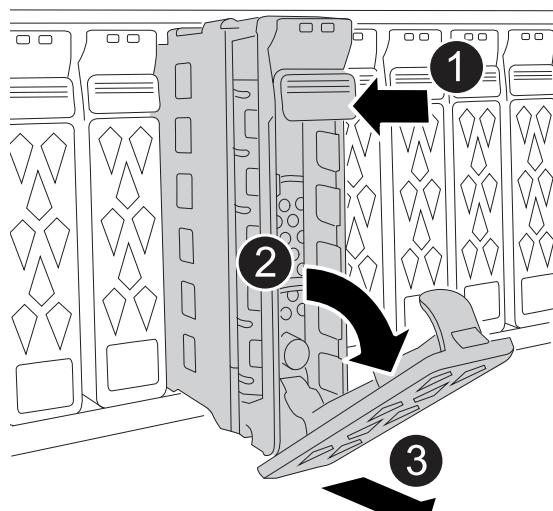
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

## Steps

1. Properly ground yourself.
2. Remove the bezel from the front of the storage system.
3. Physically identify the failed drive.
  - When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the shelf operator display panel and the failed drive illuminate.
  - The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:



1	Press the release button on the drive face to open the cam handle.
2	Rotate the cam handle downward to disengage the drive from the midplane.
3	<p>Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand.</p> <p>When removing a drive, always use two hands to support its weight.</p> <p>Because drives are fragile, minimize handling to avoid damaging them.</p>

5. Wait a minimum of 70 seconds before inserting the replacement drive.

This allows the system to recognize that a drive was removed.

6. Insert the replacement drive:

- With the cam handle in the open position, use both hands to insert the replacement drive.
- Gently push until the drive stops.
- Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

7. Verify that the drive's activity (green) LED is illuminated.

When the drive's activity LED is solid, it means that the drive has power. When the drive's activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat Steps 3 through 7.

9. Reinstall the bezel on the front of the storage system.

10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Contact technical support [NetApp Support](#) if you need the RMA number or additional help with the replacement procedure.

## Replace a fan module - ASA A20, ASA A30, and ASA A50

Replace a fan module in your ASA A20, ASA A30, or ASA A50 storage system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

### About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

### Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

#### About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, **quorum status** of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

#### Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <code>y</code> when prompted.

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

## Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

### Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

### Steps

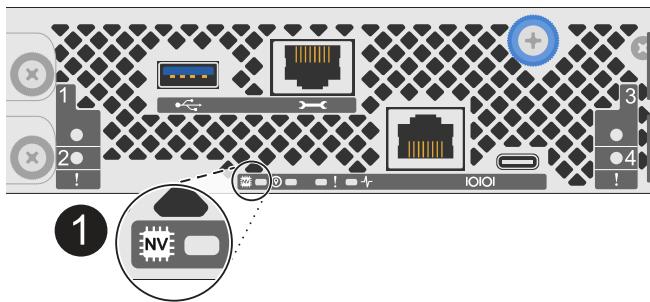
1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1

NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.

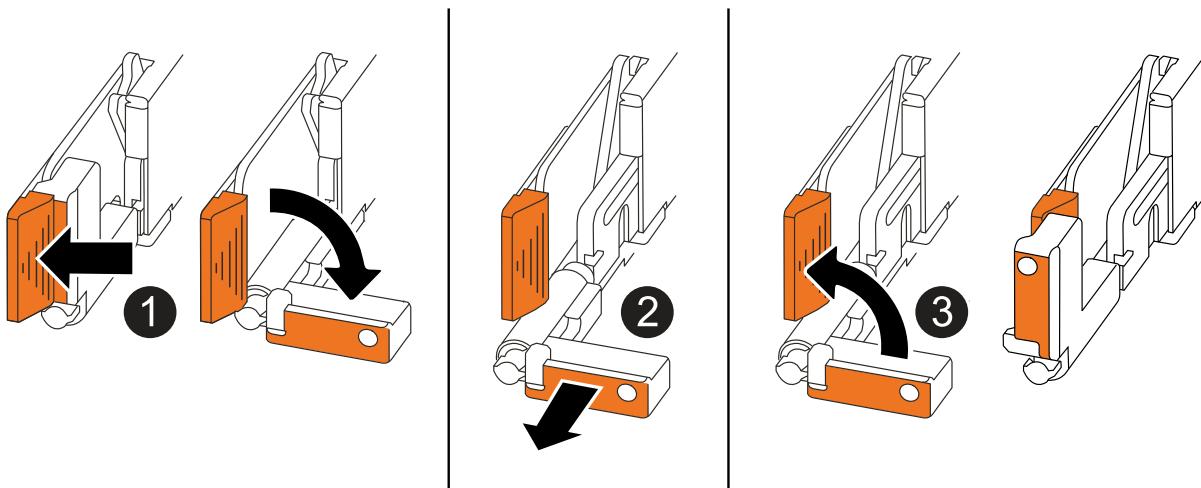
If you are disconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"> <li>1. Open the power cord retainer.</li> <li>2. Unplug the power cord from the PSU and set it aside.</li> </ol>
DC PSU	<ol style="list-style-type: none"> <li>1. Unscrew the two thumb screws on the D-SUB DC power cord connector.</li> <li>2. Unplug the power cord from the PSU and set it aside.</li> </ol>

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	<ul style="list-style-type: none"> <li>• Pull the handles towards you to unseat the controller from the midplane.</li> </ul> <p>As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling.</p> <ul style="list-style-type: none"> <li>• Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.</li> </ul>
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

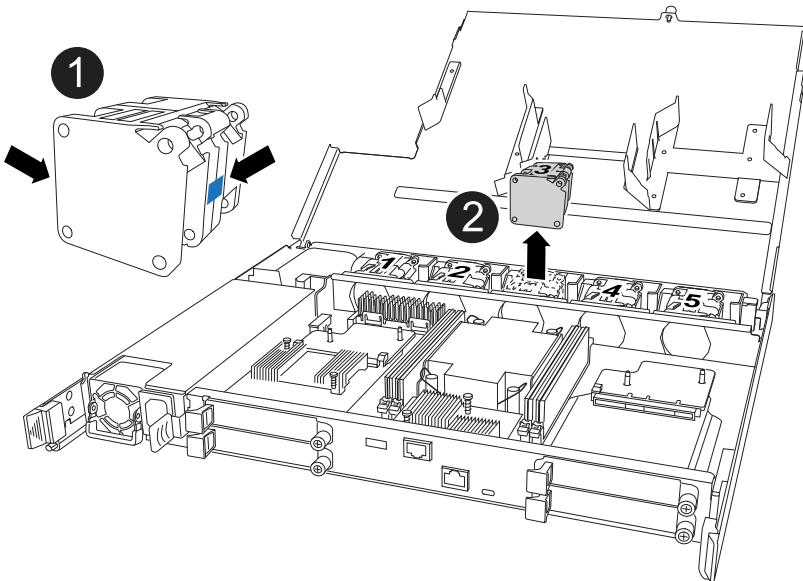
6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

### Step 3: Replace fan

To replace a fan, remove the failed fan and replace it with a new fan.

#### Steps

1. Identify the fan that you must replace by checking the console error messages.
2. Remove the failed fan:



1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

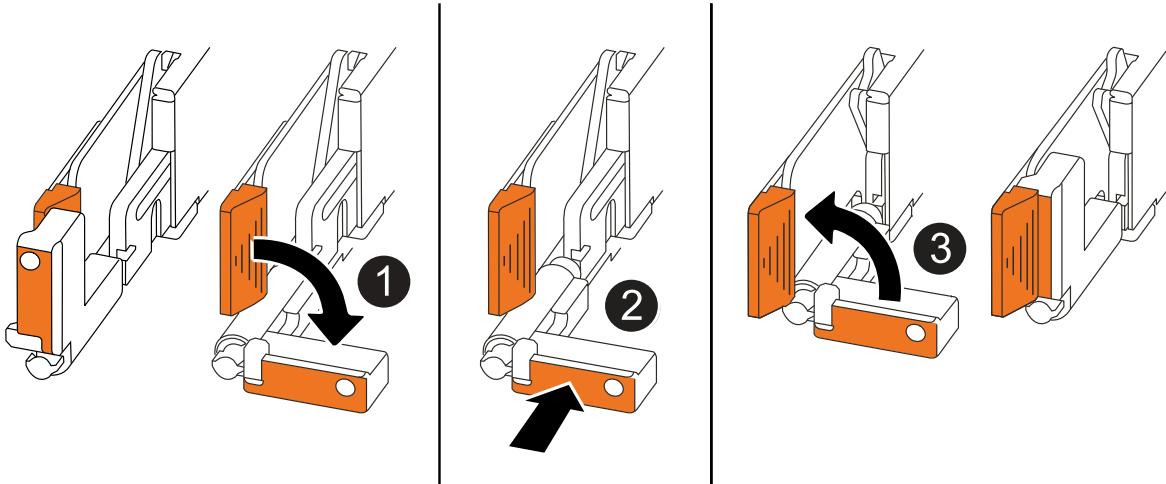
3. Insert the replacement fan by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.

### Step 4: Reinstall the controller module

Reinstall the controller into the chassis and reboot it.

#### About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

## Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.

2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.



Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:

a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

b. Rotate the controller handles up and lock in place with the tabs.



The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"> <li>1. Plug the power cord into the PSU.</li> <li>2. Secure the power cord with the power cord retainer.</li> </ol>
DC PSU	<ol style="list-style-type: none"> <li>1. Plug the D-SUB DC power cord connector into the PSU.</li> <li>2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.</li> </ol>

7. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

8. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

## Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

## I/O module

### Overview of I/O module maintenance - ASA A20, ASA A30, and ASA A50

The ASA A20, ASA A30, and ASA A50 storage systems offer flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding, hot-swapping or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your storage system with the same type of I/O module, or with a different type of I/O module. You can hot-swap a cluster and HA I/O module when your storage system meets specific requirements. You can also add an I/O module to a storage system with available slots.

- [Add an I/O module](#)

Adding additional I/O modules can improve redundancy, helping to ensure that the storage system remains operational even if one I/O module fails.

- [Hot-swap an I/O module](#)

You can hot-swap certain I/O modules for an equivalent I/O module to restore the storage system to its optimal operating state. Hot-swapping is done without having to perform a manual takeover.

To use this procedure, your storage system must be running ONTAP 9.17.1 or later and meet specific system requirements.

- [Replace an I/O module](#)

Replacing a failing I/O module can restore the storage system to its optimal operating state.

## Add an I/O module - ASA A20, ASA A30, and ASA A50

Add an I/O module to your ASA A20, ASA A30, or ASA A50 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 ASA A20, ASA A30, or ASA A50 storage system if there are available slots. If all slots are fully populated, you can replace an existing module to add a new one.

### About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the asaected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

### Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

### About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, [quorum status](#) of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

### Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- Enter *y* when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller:  <code>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</code>  The <i>-halt true</i> parameter brings you to the LOADER prompt.

### Step 2: Add the new I/O module

If the storage system has available slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

#### Before you begin

- Check the [NetApp Hardware Universe](#) to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in [NetApp Hardware Universe](#) and use the best one available for your I/O module.
- All other components in the storage system must be functioning properly; if not, contact [NetApp Support](#) before continuing with this procedure.

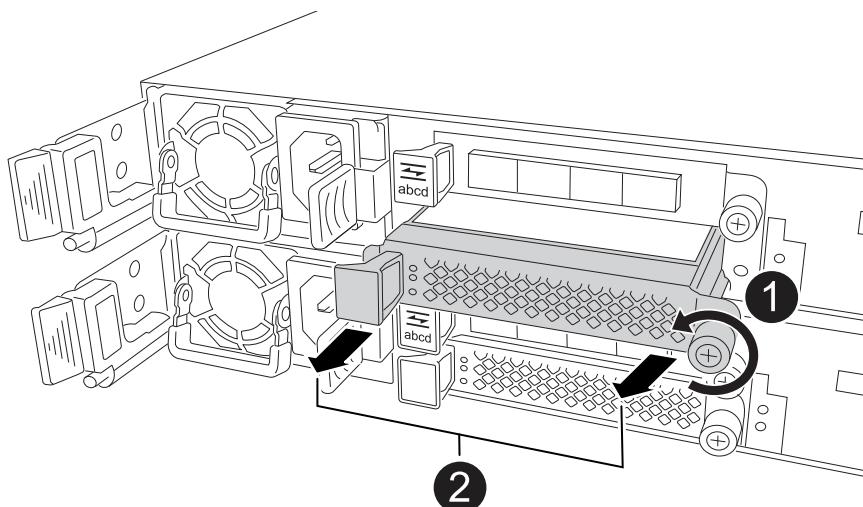
## Add I/O module to an available slot

You can add a new I/O module into a storage system with available slots.

### Steps

1. If you are not already grounded, properly ground yourself.
2. On the impaired controller, remove the I/O blanking module from the target slot.

Unused I/O slots should have blanking module installed to prevent possible thermal issues and for EMC compliance.



1	On the I/O blanking module, turn the thumbscrew counterclockwise to loosen.
2	Pull the I/O blanking module out of the controller using the tab on the left and the thumbscrew.

### 3. Install the new I/O module:

- a. Align the I/O module with the edges of the controller slot opening.
- b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O module.

- c. Turn the thumbscrew clockwise to tighten.

### 4. Cable the I/O module to the designated devices.

If you installed a storage I/O module, install and cable your NS224 shelves, as described in [Hot-add workflow](#).

### 5. Reboot the impaired controller from the LOADER prompt: `bye`

Rebooting the impaired controller also reinitializes the I/O modules and other components.

### 6. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name.
```

7. Repeat these steps to add an I/O module to the other controller.
8. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation: +

```
system node autosupport invoke -node * -type all -message MAINT=END
```

### Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

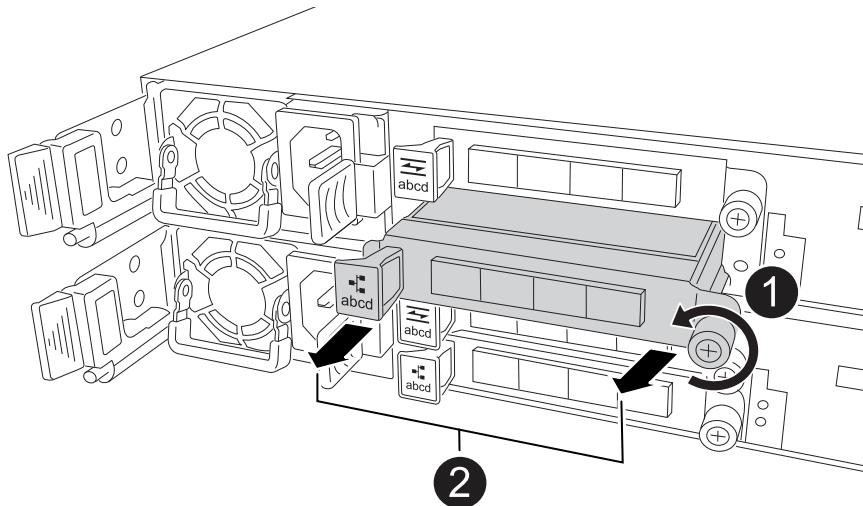
#### About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See <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. On the impaired controller, unplug any cabling on the target I/O module.
3. Remove the target I/O module from the controller:



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the new I/O module into the target slot:

- a. Align the I/O module with the edges of the slot.
- b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O Module.

- c. Turn the thumbscrew clockwise to tighten.

5. Cable the I/O module to the designated devices.

If you installed a storage I/O module, install and cable your NS224 shelves, as described in [Hot-add workflow](#).

6. Repeat the I/O module remove and install steps to add any additional I/O modules in the controller.

7. Reboot the impaired controller from the LOADER prompt:

bye

Rebooting the impaired controller also reinitializes the I/O modules and other components.

8. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

9. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

10. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

11. If you installed a NIC module, specify the usage mode for each port as *network*:

```
storage port modify -node node_name -port port_name -mode network
```

12. Repeat these steps for the other controller.

### Hot swap an I/O module - ASA A20, ASA A30, and ASA A50

You can hot swap an Ethernet I/O module in your ASA A20, ASA A30, or ASA A50 storage system if a module fails and your storage system meets all ONTAP version requirements.

To hot swap an I/O module, ensure your storage system meets the ONTAP version requirements, 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

- Hot-swapping the I/O module means that you do not have to perform a manual takeover before replacing the failed I/O module.
- Apply commands to the correct controller and I/O slot when you are hot-swapping the I/O module:
  - The *impaired controller* is the controller on which you are hot-swapping 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.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

### Step 1: Ensure the storage system meets the procedure requirements

To use this procedure, your storage system must be running ONTAP 9.17.1 or later, and your storage system must meet all requirements for the version of ONTAP your storage system is running.



If your storage system is not running ONTAP 9.17.1 or later, or does not meet all requirements for the version of ONTAP your storage system is running, you cannot use this procedure, you must use the [replace an I/O module procedure](#).

## ONTAP 9.17.1 or 9.18.1RC

- You are hot-swapping a failed cluster and HA I/O module in slot 4 with an equivalent I/O module. You cannot change the I/O module type.
- The controller with the failed cluster and HA I/O module (the impaired controller) must have already taken over the healthy partner controller. The takeover should have occurred automatically if the I/O module has failed.

For two-node clusters, the storage system cannot discern which controller has the failed I/O module, so either controller might initiate the takeover. Hot swapping is only supported when the controller with the failed I/O module (the impaired controller) has taken over the healthy controller. Hot-swapping the I/O module is the only way to recover without an outage.

You can verify that the impaired controller successfully took over the healthy controller by entering the `storage failover show` command.

If you are not sure which controller has the failed I/O module, contact [NetApp Support](#).

- Your storage system configuration must have only one cluster and HA I/O module located in slot 4, not two cluster and HA I/O modules.
- Your storage system must be a two-node (switchless or switched) cluster configuration.
- All other components in the storage system must be functioning properly; if not, contact [NetApp Support](#) before continuing with this procedure.

## ONTAP 9.18.1GA or later

- 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 controllers 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. Unplug the cables from the failed I/O module.

Make sure to label the cables so you can reconnect them to the same ports later in this procedure.

 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. As needed for the version of ONTAP your storage system is running and the state of the controllers, disable automatic giveback:

ONTAP version	If...	Then...
9.17.1 or 9.18.1RC	If the impaired controller took over the healthy controller automatically	<p>Disable automatic giveback:</p> <ol style="list-style-type: none"><li>1. Enter the following command from the console of the impaired controller</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>
9.18.1GA or later	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>

ONTAP version	If...	Then...
9.18.1GA or later	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:

- Enter the following command:

```
system controller slot module remove -node impaired_node_name -slot
slot_number
```

- Enter *y* when you see the prompt *Do you want to continue?*

For example, the following command prepares the failed module in slot 4 on node 2 (the impaired controller) for removal, and displays a message that it is safe to remove:

```
node2::> system controller slot module remove -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered off for removal.

Do you want to continue? {y|n}: y

The module has been successfully removed from service and powered
off. It can now be safely removed.
```

6. Verify the failed I/O module is powered off:

```
system controller slot module show
```

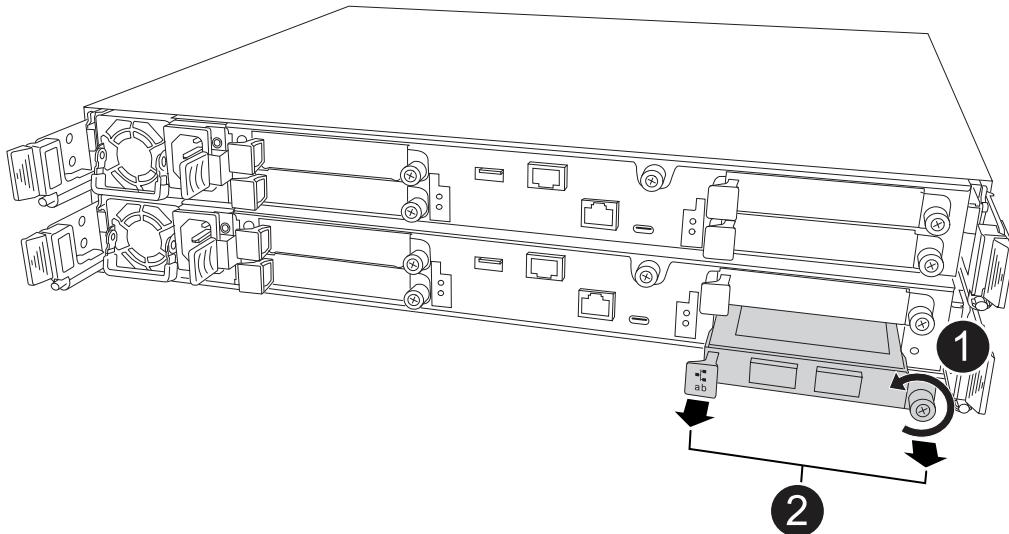
The output should show *powered-off* in the *status* column for the failed module and its slot number.

### Step 3: Hot swap the failed I/O module

Hot swap the failed I/O module with an equivalent I/O module:

#### Steps

- If you are not already grounded, properly ground yourself.
- Remove the failed I/O module from the impaired controller:



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew on the right.

3. Install the replacement I/O module:

- Align the I/O module with the edges of the slot.
- Gently push the I/O module all the way into the slot, making sure to properly seat the I/O module into the connector.

You can use the tab on the left and the thumbscrew on the right to push in the I/O module.

- Turn the thumbscrew clockwise to tighten.

4. Cable the replacement I/O module.

**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:

- Enter the following command:

```
system controller slot module insert -node impaired_node_name -slot
slot_number
```

- 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 4 on node 2 (the impaired controller) online, and displays a message that the process was successful:

```
node2::> system controller slot module insert -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered on and initialized.

Do you want to continue? {y|n}: `y`

The module has been successfully powered on, initialized and placed
into service.
```

2. Verify that each port on the I/O module successfully initialized:

- 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 indicating each port on the I/O module initiated successfully.

For example, the following output shows initialization succeeded for I/O ports e4b and e4a:

```
node2::> event log show -event *hotplug.init*

Time           Node           Severity      Event
-----
-----
7/11/2025 16:04:06  node2      NOTICE       hotplug.init.success:
Initialization of ports "e4b" in slot 4 succeeded

7/11/2025 16:04:06  node2      NOTICE       hotplug.init.success:
Initialization of ports "e4a" in slot 4 succeeded

2 entries were displayed.
```

- If the port initialization fails, review the EMS log for the next steps to take.
- 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 a I/O module in slot 4:

```
node2::> system controller config show -node local -slot 4

Node: node2
Sub- Device/
Slot slot Information
-----
4      - Dual 40G/100G Ethernet Controller CX6-DX
          e4a MAC Address: d0:39:ea:59:69:74 (auto-100g_cr4-fd-
up)
          QSFP Vendor:      CISCO-BIZLINK
          QSFP Part Number: L45593-D218-D10
          QSFP Serial Number: LCC2807GJFM-B
          e4b MAC Address: d0:39:ea:59:69:75 (auto-100g_cr4-fd-
up)
          QSFP Vendor:      CISCO-BIZLINK
          QSFP Part Number: L45593-D218-D10
          QSFP Serial Number: LCC2809G26F-A
          Device Type:      CX6-DX PSID(NAP0000000027)
          Firmware Version: 22.44.1700
          Part Number:      111-05341
          Hardware Revision: 20
          Serial Number:    032403001370
```

#### Step 5: Restore the storage system to normal operation

Restore your storage system to normal operation by giving back storage to the 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:

ONTAP version	If...	Then...
9.17.1 or 9.18.1RC	If the impaired controller took over the healthy controller automatically	<ol style="list-style-type: none"> <li>1. Return the healthy controller to normal operation by giving back its storage:  <code>storage failover giveback -ofnode <i>healthy_node_name</i></code></li> <li>2. Restore automatic giveback from the console of the impaired controller:  <code>storage failover modify -node local -auto-giveback true</code></li> </ol>
9.18.1GA or later	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 <i>controller that was taken over_name</i></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>
9.18.1GA or later	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 server and ports: `network interface show -is-home false`

If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

3. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

#### Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

#### Replace an I/O module - ASA A20, ASA A30, and ASA A50

Replace an I/O module in your ASA A20, ASA A30, or ASA A50 storage system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

Use this procedure to replace a failed I/O module.

## Before you begin

All other components in the storage system must be functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

## About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

### Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

## About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, [quorum status](#) of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

## Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

## Step 2: Replace a failed I/O module

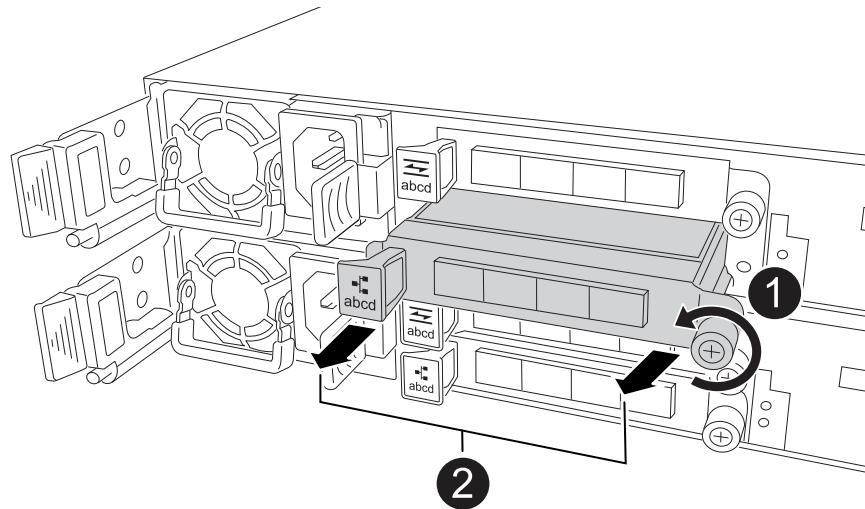
To replace a failed I/O module, locate it in the controller and follow the specific sequence of steps.

### Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug cabling from the failed I/O module.

Make sure to label the cables so that you know where they came from.

3. Remove the failed I/O module from the controller:



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the replacement I/O module into the target slot:

- a. Align the I/O module with the edges of the slot.

- b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O Module.

- c. Turn the thumbscrew clockwise to tighten.

## 5. Cable the I/O module.

### Step 3: Reboot the controller

After you replace an I/O module, you must reboot the controller.

#### Steps

1. Reboot the controller from the LOADER prompt: `bye`

Rebooting the impaired controller also reinitializes the I/O modules and other components.

2. Return the node to normal operation: `storage failover giveback -ofnode impaired_node_name`
3. Restore automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback true`

### Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

## Replace the NV battery - ASA A20, ASA A30, and ASA A50

Replace the NV battery in your ASA A20, ASA A30, or ASA A50 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.

To replace the NV battery, you must remove the controller, remove the faulty battery, install the replacement battery, and then reinstall the controller.

#### Before you begin

All other components in the storage system must be functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

#### About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

## Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

### About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, `quorum status` of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

### Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <code>y</code> when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller:  <code>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</code>  The <code>-halt true</code> parameter brings you to the LOADER prompt.

## Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

### Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

### Steps

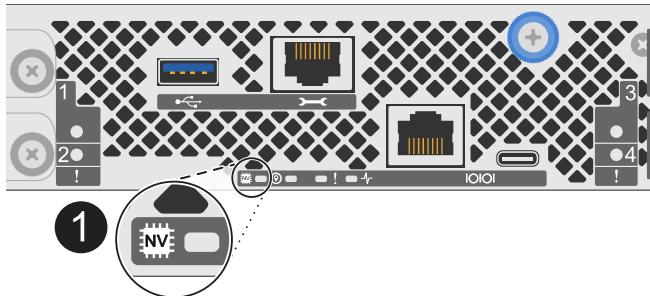
1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.

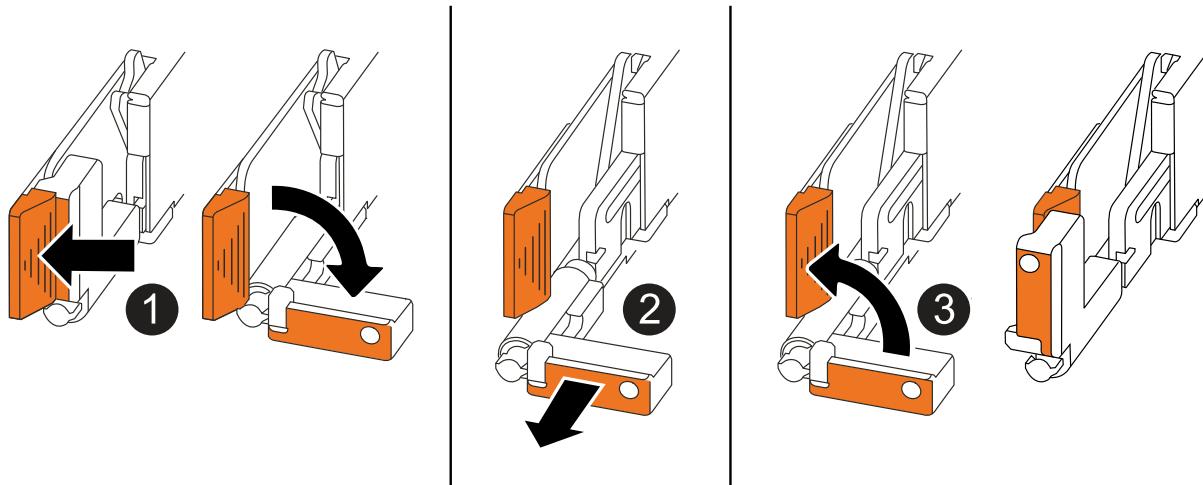
If you are disconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"><li>1. Open the power cord retainer.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>
DC PSU	<ol style="list-style-type: none"><li>1. Unscrew the two thumb screws on the D-SUB DC power cord connector.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	<ul style="list-style-type: none"><li>Pull the handles towards you to unseat the controller from the midplane.</li><li>As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling.</li><li>Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.</li></ul>
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

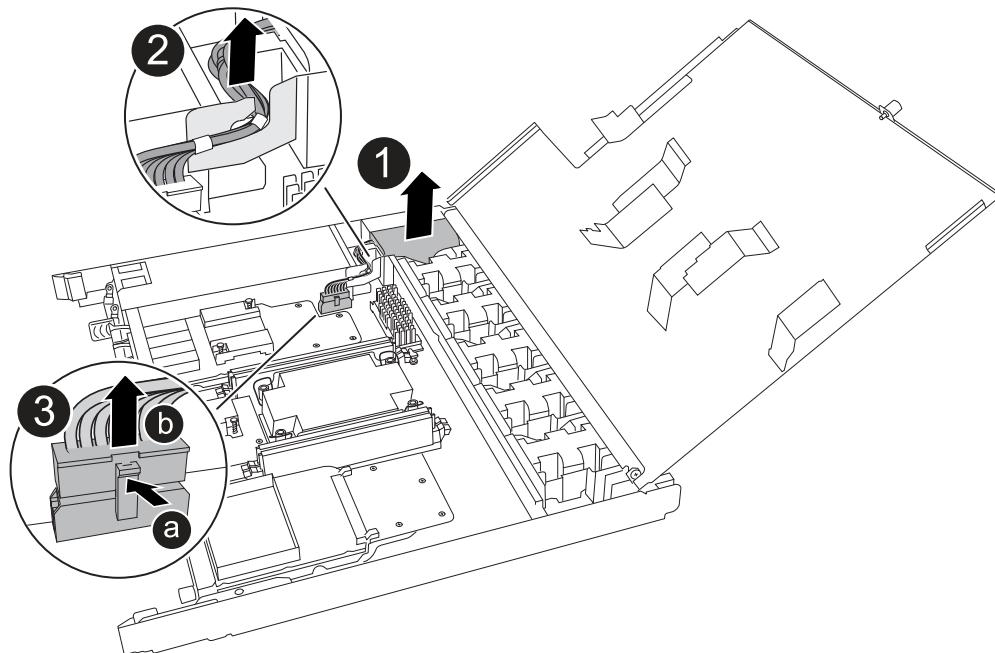
6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

### Step 3: Replace the NV battery

Remove the failed NV battery from the controller and install the replacement NV battery.

#### Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the NV battery.
3. Remove the NV battery:



1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	<ol style="list-style-type: none"> <li>Push in and hold the tab on the connector.</li> <li>Pull the connector up and out of the socket.</li> </ol> <p>As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.</p>

4. Install the replacement NV battery:

- Remove the replacement battery from its package.
- Plug the wiring connector into its socket.
- Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
- Place the NV battery into its compartment.

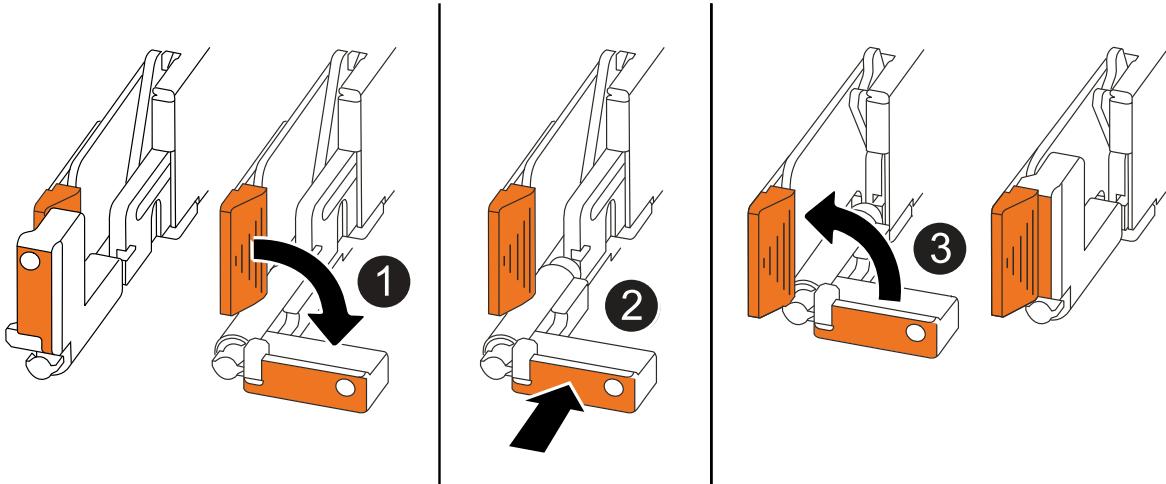
The NV battery should sit flush in its compartment.

#### Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

#### About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

## Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.

2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.



Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:

a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

b. Rotate the controller handles up and lock in place with the tabs.



The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"> <li>1. Plug the power cord into the PSU.</li> <li>2. Secure the power cord with the power cord retainer.</li> </ol>
DC PSU	<ol style="list-style-type: none"> <li>1. Plug the D-SUB DC power cord connector into the PSU.</li> <li>2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.</li> </ol>

7. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

8. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

## Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

## Hot-swap a power supply - ASA A20, ASA A30, and ASA A50

Replace an AC or DC power supply unit (PSU) in your ASA A20, ASA A30, or ASA A50 storage system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cord, replacing the faulty PSU, and then reconnecting it to the power source.

### About this task

- This procedure is written for replacing one PSU at a time.

The PSUs are redundant and hot-swappable. You do not have to shut down the controller to replace a PSU.

- **IMPORTANT:** Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.
- Use the appropriate procedure for your type of PSU: AC or DC.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected

storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

## Option 1: Hot-swap an AC PSU

To replace an AC PSU, complete the following steps.

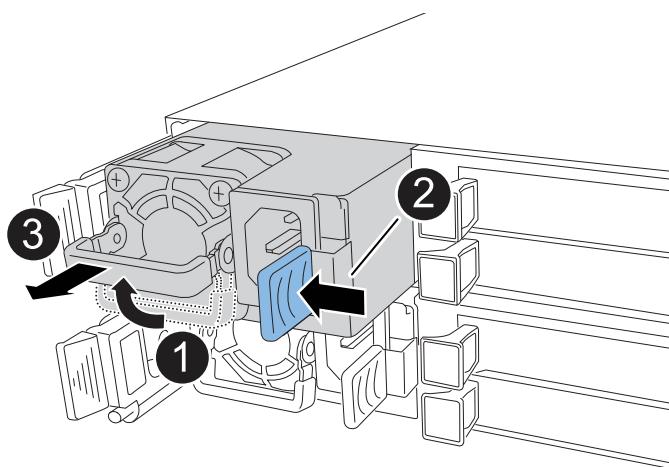
### Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the power cord from the PSU by opening the power cord retainer, and then unplug the power cord from the PSU.



PSUs do not have a power switch.

4. Remove the PSU:



1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	<p>Pull the PSU out of the controller while using your other hand to support its weight.</p> <p> The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.</p>

5. Install the replacement PSU:

- a. Using both hands, support and align the edges of the PSU with the opening in the controller.
- b. Gently push the PSU into the controller until the locking tab clicks into place.

A PSU will only properly engage with the internal connector and lock in place one way.



To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

- c. Rotate the handle down, so it is out of the way of normal operations.
6. Reconnect the power cord to the PSU and secure the power cord with the power cord retainer.

Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

### Option 2: Hot-swap a DC PSU

To replace a DC PSU, complete the following steps.

#### Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the PSU:

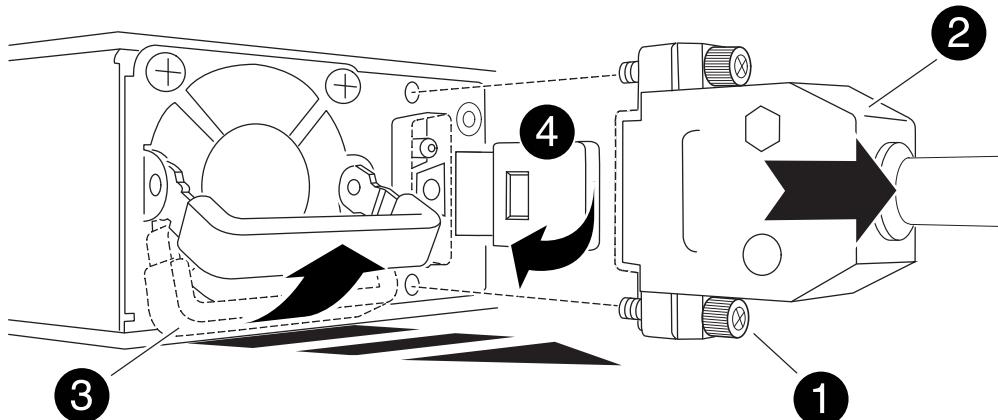


PSUs do not have a power switch.

- a. Unscrew the two thumb screws on the D-SUB DC power cord connector.
- The illustration and table in step 4 shows the two thumb screws (item #1) and the D-SUB DC power cord connector (item #2).
- b. Unplug the cord from the PSU and set it aside.
4. Remove the PSU:
  - a. Rotate the handle up, to its horizontal position, and then grasp it.
  - b. With your thumb, press the terracotta tab to release the locking mechanism.
  - c. Pull the PSU out of the controller while using your other hand to support its weight.



The PSU is short. Always use two hands to support it when removing it from the controller so that it does not swing free from the controller and injure you.



1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

5. Insert the replacement PSU:

- Using both hands, support and align the edges of the PSU with the opening in the controller.
- Gently slide the PSU into the controller until the locking tab clicks into place.

A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.



To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

- Rotate the handle down, so it is out of the way of normal operations.

6. Reconnect the D-SUB DC power cord:

Once power is restored to the PSU, the status LED should be green.

- Plug the D-SUB DC power cord connector into the PSU.
- Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

## Replace the real-time clock battery - ASA A20, ASA A30, and ASA A50

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your ASA A20, ASA A30, or ASA A50 storage system to ensure that services and applications relying on accurate time synchronization remain operational.

You replace the real-time clock (RTC) battery in the controller so that your storage system's services and applications that depend on accurate time synchronization continue to function.

### Before you begin

All other components in the storage system must be functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

### About this task

- You can use this procedure with all versions of ONTAP supported by your storage system.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

### Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

#### About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, `quorum status` of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

#### Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <code>y</code> when prompted.

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

## Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

### Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

### Steps

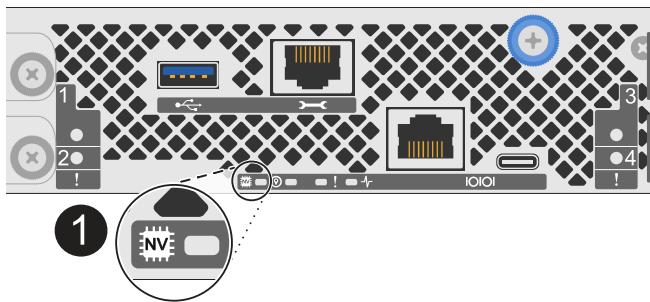
1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1

NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.

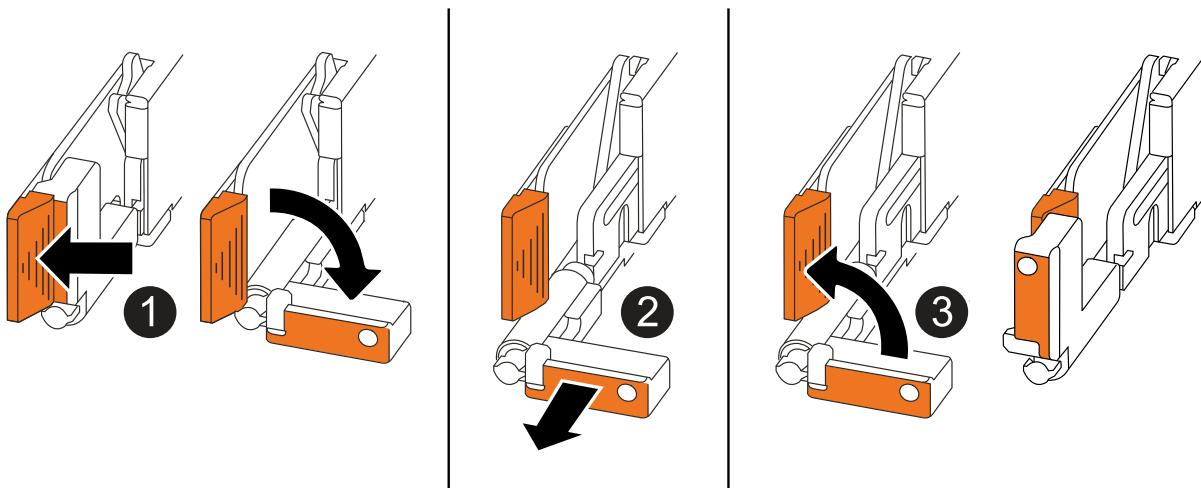
If you are disconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"> <li>1. Open the power cord retainer.</li> <li>2. Unplug the power cord from the PSU and set it aside.</li> </ol>
DC PSU	<ol style="list-style-type: none"> <li>1. Unscrew the two thumb screws on the D-SUB DC power cord connector.</li> <li>2. Unplug the power cord from the PSU and set it aside.</li> </ol>

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	<ul style="list-style-type: none"> <li>• Pull the handles towards you to unseat the controller from the midplane.</li> </ul> <p>As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling.</p> <ul style="list-style-type: none"> <li>• Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.</li> </ul>
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

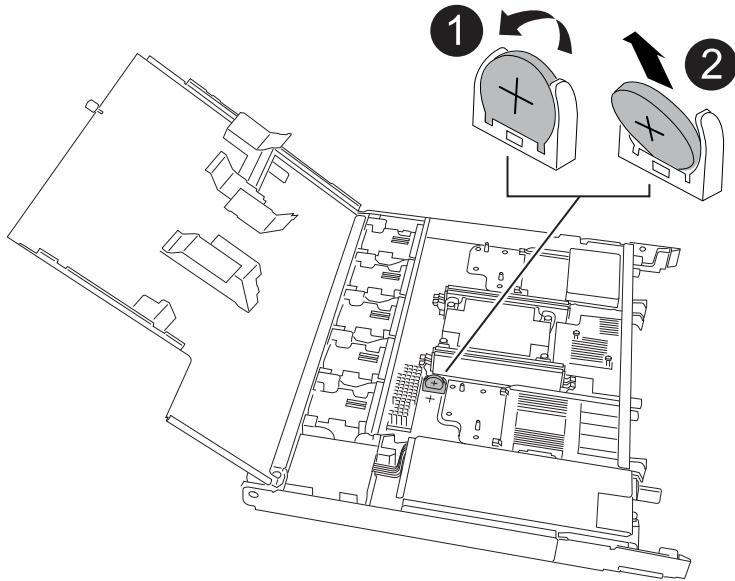
6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

### Step 3: Replace the RTC battery

Remove the failed RTC battery and install the replacement RTC battery.

#### Steps

1. Locate the RTC battery.
2. Remove the RTC battery:



1	Gently rotate the RTC battery at an angle away from its holder.
2	Lift the RTC battery out of its holder.

3. Install the replacement RTC battery:

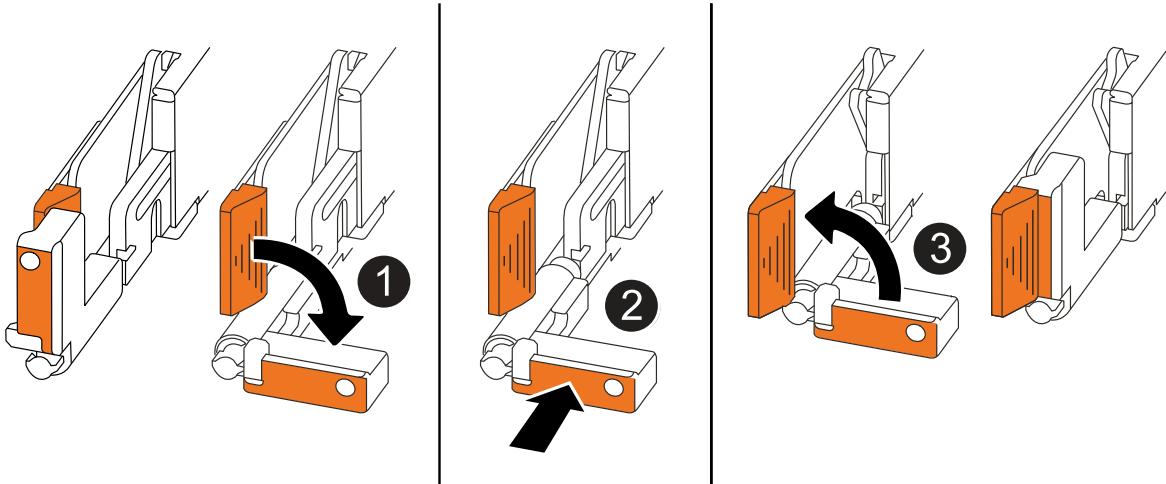
- a. Remove the replacement battery from the antistatic shipping bag.
- b. Position the battery so that the plus sign on the battery faces out to correspond with the plus sign on the motherboard.
- c. Insert the battery into the holder at an angle, and then push it into an upright position so it is fully seated in the holder.
- d. Visually inspect the battery to make sure that it is completely seated in its holder and that the polarity is correct.

### Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

#### About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

## Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.

2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.



Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:

a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

b. Rotate the controller handles up and lock in place with the tabs.



The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"> <li>1. Plug the power cord into the PSU.</li> <li>2. Secure the power cord with the power cord retainer.</li> </ol>
DC PSU	<ol style="list-style-type: none"> <li>1. Plug the D-SUB DC power cord connector into the PSU.</li> <li>2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.</li> </ol>

7. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

8. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

## Step 5: Reset the time and date on the controller

After replacing the RTC battery, inserting the controller and powering on first BIOS reset, you will see the following error messages:



RTC date/time error. Reset date/time to default

RTC power failure error

These messages are expected and you can continue with this procedure.

1. On the healthy controller, check the date and time with the `cluster date show` command.
 

If your storage system stops at the boot menu, select the option for `Reboot node` and respond `y` when prompted, then boot to `LOADER` by pressing `Ctrl-C`.
2. On the impaired controller, at the `LOADER` prompt, check the time and date: `cluster date show`
  - a. If necessary, modify the date: `set date mm/dd/yyyy`
  - b. If necessary, set the time, in GMT: `set time hh:mm:ss`
  - c. Confirm the date and time.
3. At the `LOADER` prompt, enter `bye` to reinitialize the I/O modules, other components, and let the controller reboot.
4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`

5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

## 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.

# ASA C30 systems

## Overview of hardware maintenance - ASA C30

Maintain the hardware of your ASA C30 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 procedures in this section assume that the ASA C30 storage system has already been deployed as a storage node in the ONTAP environment.

### System components

For the ASA C30 storage system, you can perform maintenance procedures on the following components.

<a href="#">Boot media - automated recovery</a>	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media.
<a href="#">Chassis</a>	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
<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="#">Drive</a>	A drive is a device that provides the physical storage needed for data.
<a href="#">Fan</a>	A fan cools the controller and drives.
<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.

NV battery	The non-volatile memory (NV) battery is responsible for providing power to the NVMEM components while data in-flight is being destaged to flash memory after a power loss.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.

## Boot media

### Boot media replacement workflow - ASA C30

Get started with replacing the boot media in your ASA C30 storage system by reviewing the replacement requirements, shutting down the impaired controller, replacing the boot media, restoring the image on the boot media, and verifying the system functionality.

1

#### Review the boot media requirements

Review the requirements for boot media replacement.

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 boot media

Remove the failed boot media from the impaired controller and install the replacement boot media.

4

#### Restore the image on the boot media

Restore the ONTAP image from the healthy 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 - ASA C30

Before replacing the boot media in your ASA C30 storage system, ensure you meet the necessary requirements and considerations for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0M (wrench) port on the impaired controller is working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

## What's next

After you've reviewed the boot media requirements, you [shut down the controller](#).

### Shut down the controller to replace the boot media - ASA C30

Shut down the impaired controller in your ASA C30 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, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

#### About this task

- If you have a SAN system, you must have checked event messages (cluster kernel-service show) for the impaired controller SCSI blade. The cluster kernel-service show command (from priv advanced mode) displays the node name, [quorum status](#) of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

#### Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

## 2. Disable automatic giveback:

- Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- Enter *y* when you see the prompt *Do you want to disable auto-giveback?*

## 3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> The <i>-halt true</i> parameter brings you to the LOADER prompt.

## What's next

After you shut down the impaired controller, you [replace the boot media](#).

## Replace the boot media - ASA C30

The boot media in your ASA C30 storage system stores essential firmware and configuration data. The replacement process involves removing the controller module, removing the impaired boot media, installing the replacement boot media, and then transferring the ONTAP image to the replacement boot media.

### About this task

If needed, you can turn on the platform chassis location (blue) LEDs to aid in physically locating the affected platform. Log into the BMC using SSH and enter the `system location-led on` command.

A platform chassis has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

### Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

## Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

## Steps

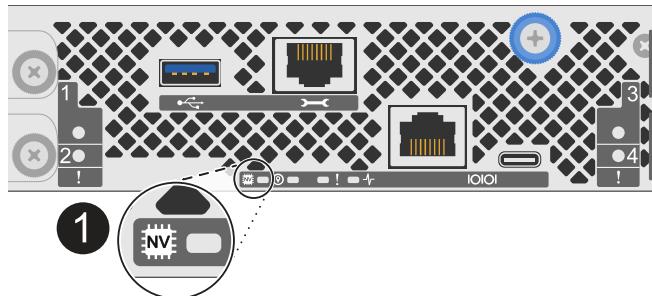
1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.

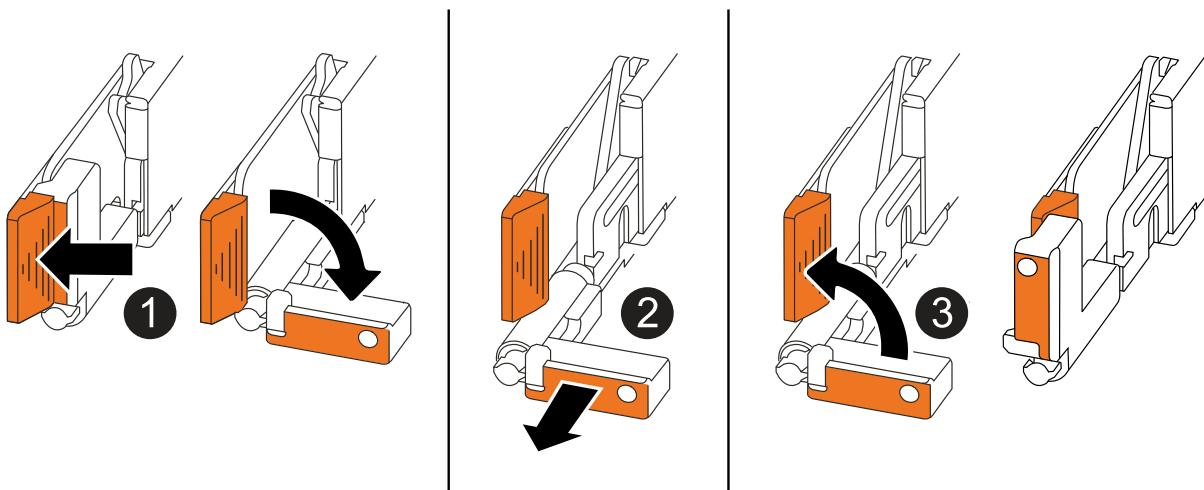
If you are disconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"><li>1. Open the power cord retainer.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>
DC PSU	<ol style="list-style-type: none"><li>1. Unscrew the two thumb screws on the D-SUB DC power cord connector.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



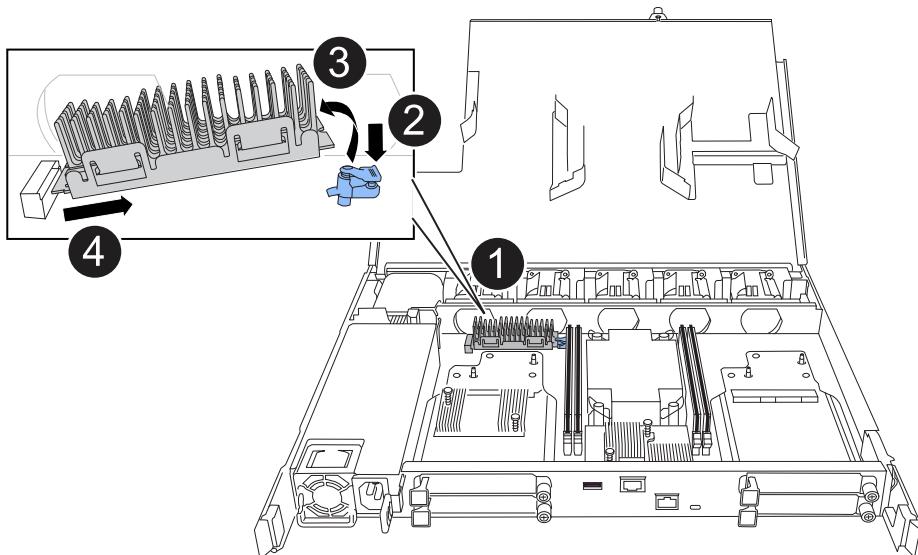
1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	<ul style="list-style-type: none"> <li>Pull the handles towards you to unseat the controller from the midplane.</li> </ul> <p>As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling.</p> <ul style="list-style-type: none"> <li>Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.</li> </ul>
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Place the controller on an anti-static mat.
7. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

#### Step 2: Replace the boot media

To replace the boot media, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Remove the boot media:



1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

3. Install the replacement boot media:

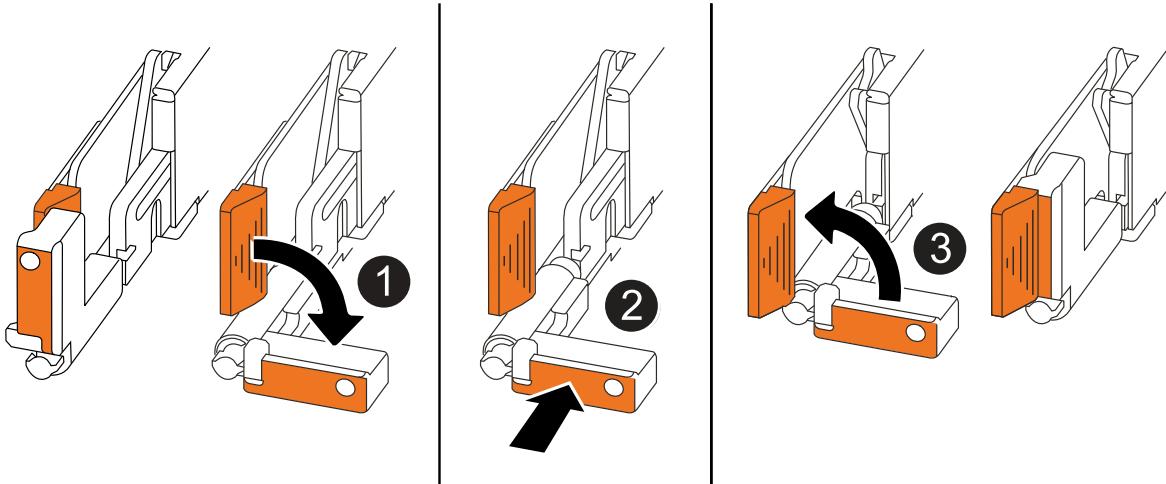
- Remove the boot media from its package.
- Slide the socket end of the boot media into its socket.
- At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

**Step 3: Reinstall the controller**

Reinstall the controller into the chassis and reboot it.

**About this task**

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

## Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



Do not completely insert the controller in the chassis until instructed to do so later in this procedure.

3. Reconnect the cables to the controller; however, do not plug in the power cord to the power supply (PSU) at this time.



Make sure that the console cable is connected to the controller because you want to catch and log the boot sequence later in the boot media replacement procedure when you fully seat the controller in the chassis and it begins to boot.

4. Fully seat the controller in the chassis:

- a. Firmly push on the handles until the controller meets the midplane and is fully seated.

Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.



The controller boots to the LOADER prompt when fully seated in the chassis. It gets its power from the partner controller.

b. Rotate the controller handles up and lock in place with the tabs.

5. Reconnect the power cord to the PSU on the impaired controller.

Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"> <li>1. Plug the power cord into the PSU.</li> <li>2. Secure the power cord with the power cord retainer.</li> </ol>
DC PSU	<ol style="list-style-type: none"> <li>1. Plug the D-SUB DC power cord connector into the PSU.</li> <li>2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.</li> </ol>

### What's next

After physically replacing the impaired boot media, you [restore the ONTAP image from the partner node](#).

### Restore the ONTAP image on the boot media - ASA C30

After installing the new boot media device in your ASA C30 storage system, you can start the automated boot media recovery process to restore the configuration from the healthy node.

During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

### Before you begin

- Determine your key manager type:
  - Onboard Key Manager (OKM): Requires cluster-wide passphrase and backup data
  - External Key Manager (EKM): Requires the following files from the partner node:
    - /cfcard/kmip/servers.cfg
    - /cfcard/kmip/certs/client.crt
    - /cfcard/kmip/certs/client.key
    - /cfcard/kmip/certs/CA.pem

### Steps

1. From the LOADER prompt, start the boot media recovery process:

```
boot_recovery -partner
```

The screen displays the following message:

```
Starting boot media recovery (BMR) process. Press Ctrl-C to abort...
```

2. Monitor the boot media install recovery process.

The process completes and displays the Installation complete message.

3. The system checks for encryption and displays one of the following messages:

If you see this message...	Do this...
key manager is not configured. Exiting.	<p>Encryption is not installed on the system.</p> <ol style="list-style-type: none"><li>Wait for the login prompt to display.</li><li>Log into the node and give back the storage: <pre>storage failover giveback -ofnode impaired_node_name</pre></li><li>Go to <a href="#">re-enabling automatic giveback</a> if it was disabled.</li></ol>
key manager is configured.	Encryption is installed. Go to <a href="#">restoring the key manager</a> .



If the system cannot identify the key manager configuration, it displays an error message and prompts you to confirm whether key manager is configured and which type (onboard or external). Answer the prompts to proceed.

4. Restore the key manager using the appropriate procedure for your configuration:

## Onboard Key Manager (OKM)

The system displays the following message and begins running BootMenu Option 10:

```
key manager is configured.  
Entering Bootmenu Option 10...
```

```
This option must be used only in disaster recovery procedures. Are  
you sure? (y or n):
```

- a. Enter **y** at the prompt to confirm you want to start the OKM recovery process.
- b. Enter the passphrase for onboard key management when prompted.
- c. Enter the passphrase again when prompted to confirm.
- d. Enter the backup data for onboard key manager when prompted.

### Show example of passphrase and backup data prompts

```
Enter the passphrase for onboard key management:  
-----BEGIN PASSPHRASE-----  
<passphrase_value>  
-----END PASSPHRASE-----  
Enter the passphrase again to confirm:  
-----BEGIN PASSPHRASE-----  
<passphrase_value>  
-----END PASSPHRASE-----  
Enter the backup data:  
-----BEGIN BACKUP-----  
<passphrase_value>  
-----END BACKUP-----
```

- e. Monitor the recovery process as it restores the appropriate files from the partner node.

When the recovery process is complete, the node reboots. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....  
Setting recovery material for the onboard key manager  
Recovery secrets set successfully  
Trying to delete any existing km_onboard.keydb file.  
  
Successfully recovered keymanager secrets.
```

- f. After the node reboots, verify that the system is back online and operational.

g. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

h. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster:

```
security key-manager onboard sync
```

Go to [re-enabling automatic giveback](#) if it was disabled.

### External Key Manager (EKM)

The system displays the following message and begins running BootMenu Option 11:

```
key manager is configured.  
Entering Bootmenu Option 11...
```

a. Enter the EKM configuration settings when prompted:

i. Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file:

#### Show example of client certificate contents

```
-----BEGIN CERTIFICATE-----  
<certificate_value>  
-----END CERTIFICATE-----
```

ii. Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file:

#### Show example of client key file contents

```
-----BEGIN RSA PRIVATE KEY-----  
<key_value>  
-----END RSA PRIVATE KEY-----
```

iii. Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file:

#### Show example of KMIP server file contents

```
-----BEGIN CERTIFICATE-----  
<KMIP_certificate_CA_value>  
-----END CERTIFICATE-----
```

iv. Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file:

**Show example of server configuration file contents**

```
xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx
xxx.xxx.xxx.xxx:5696.port=5696
xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem
xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4
1xxx.xxx.xxx.xxx:5696.timeout=25
xxx.xxx.xxx.xxx:5696.nbio=1
xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt
xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key
xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:
!RC2:!RC4:!SEED:!eNULL:!aNULL"
xxx.xxx.xxx.xxx:5696.verify=true
xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>
```

v. If prompted, enter the ONTAP Cluster UUID from the partner node. You can check the cluster UUID from the partner node using the `cluster identify show` command.

**Show example of ONTAP Cluster UUID prompt**

```
Notice: bootarg.mgwd.cluster_uuid is not set or is empty.
Do you know the ONTAP Cluster UUID? {y/n} y
Enter the ONTAP Cluster UUID: <cluster_uuid_value>
```

```
System is ready to utilize external key manager(s).
```

vi. If prompted, enter the temporary network interface and settings for the node:

- The IP address for the port
- The netmask for the port
- The IP address of the default gateway

### Show example of temporary network setting prompts

In order to recover key information, a temporary network interface needs to be configured.

```
Select the network port you want to use (for example,  
'e0a')  
e0M
```

```
Enter the IP address for port : xxx.xxx.xxx.xxx  
Enter the netmask for port : xxx.xxx.xxx.xxx  
Enter IP address of default gateway: xxx.xxx.xxx.xxx  
Trying to recover keys from key servers....  
[discover_versions]  
[status=SUCCESS reason= message=]
```

#### b. Verify the key restoration status:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored. The process restores the appropriate files from the partner node and reboots the node. Proceed to the next step.
- If the key is not successfully restored, the system halts and displays error and warning messages. Rerun the recovery process from the **LOADER** prompt: `boot_recovery -partner`

### Show example of key recovery error and warning messages

```
ERROR: kmip_init: halting this system with encrypted
mroot...
WARNING: kmip_init: authentication keys might not be
available.
*****
*          A T T E N T I O N          *
*
*      System cannot connect to key managers.      *
*
*****
ERROR: kmip_init: halting this system with encrypted
mroot...
.
Terminated

Uptime: 11m32s
System halting...

LOADER-B>
```

- c. After the node reboots, verify that the system is back online and operational.
- d. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

Go to [re-enabling automatic giveback](#) if it was disabled.

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

#### What's next

After you've restored the ONTAP image and the node is up and serving data, you [return the failed part to NetApp](#).

#### Return the failed part to NetApp - ASA C30

When a component in your ASA C30 storage system fails, return the failed part to NetApp. See the [Part Return and Replacements](#) page for further information.

# Chassis

## Chassis replacement workflow - ASA C30

Replacing the chassis in your ASA C30 storage system consists of reviewing the replacement requirements, shutting down the controllers, replacing the chassis, and verifying system operations.

1

### Review the chassis replace requirements

Review the requirements to replace the chassis.

2

### Shut down the controllers

Shut down the controllers so you can perform maintenance on the chassis.

3

### Replace the chassis

Replace the chassis by moving the drives and any drive blanks, controllers (with the power supplies), and bezel from the impaired chassis to the new chassis, and swapping out the impaired chassis with the new chassis of the same model as the impaired chassis.

4

### Complete chassis replacement

Verify the HA state of the chassis and return the failed part to NetApp.

## Requirements to replace the chassis - ASA C30

Before replacing the chassis of your ASA C30 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement chassis, and the necessary tools.

Review the following requirements and considerations.

### Requirements

- The replacement chassis must be the same model as the impaired chassis. This procedure is for a like-for-like replacement, not for an upgrade.
- All other components in the storage system must be functioning properly; if not, contact [NetApp Support](#) before continuing with this procedure.

### Considerations

- The chassis replacement procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.
- You can use the chassis replacement procedure with all versions of ONTAP supported by your storage system.

- The chassis replacement procedure is written with the assumption that you are moving the bezel, drives, any drive blanks, and controllers to the new chassis.

## What's next?

After you've reviewed the requirements to replace the chassis, you need to [shut down the controllers](#)

### Shut down the controllers - ASA C30

Shut down the controllers in your ASA C30 storage system to prevent data loss and ensure system stability when replacing the chassis.

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see [Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base](#).

#### Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessibility for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional [system health checks](#).
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any [Active IQ Wellness Alerts and Risks](#).  
Make note of any faults presently on the system, such as LEDs on the system components.

#### Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

```
system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
```

5. Identify the SP/BMC address of all cluster nodes:

```
system service-processor show -node * -fields address
```

6. Exit the cluster shell:

```
exit
```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown  
true -ignore-quorum-warnings true -inhibit-takeover true
```



For clusters using SnapMirror synchronous operating in StrictSync mode: system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict -sync-warnings true

9. Enter **y** for each controller in the cluster when you see:

Warning: Are you sure you want to halt node <node\_name>? {y|n}:

10. Wait for each controller to halt and display the LOADER prompt.

**What's next?**

After you've shut down the controllers, you need to [replace the chassis](#).

### Replace the chassis - ASA C30

Replace the chassis of your ASA C30 storage system when a hardware failure requires it. The replacement process involves removing the controllers, removing the drives, installing the replacement chassis, and reinstalling the chassis components.

#### About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

#### Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

#### Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

#### Steps

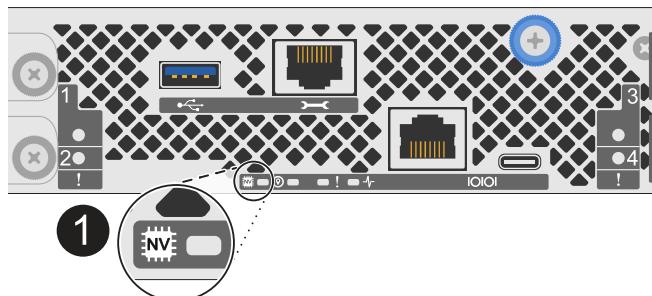
1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



**1** NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.

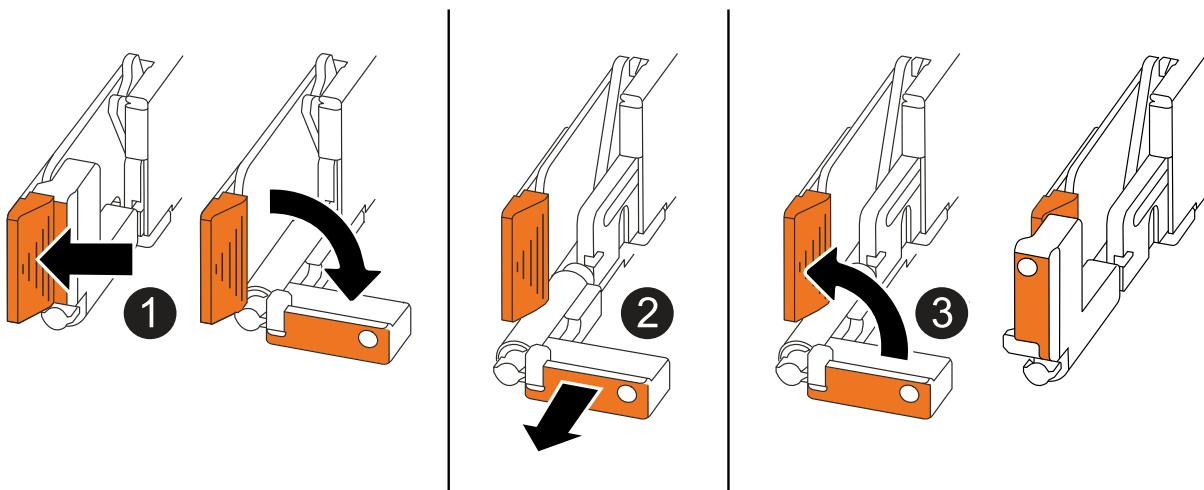
If you are disconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"><li>1. Open the power cord retainer.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>
DC PSU	<ol style="list-style-type: none"><li>1. Unscrew the two thumb screws on the D-SUB DC power cord connector.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	<ul style="list-style-type: none"> <li>Pull the handles towards you to unseat the controller from the midplane.</li> </ul> <p>As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling.</p> <ul style="list-style-type: none"> <li>Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.</li> </ul>
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Repeat these steps for the other controller in the chassis.

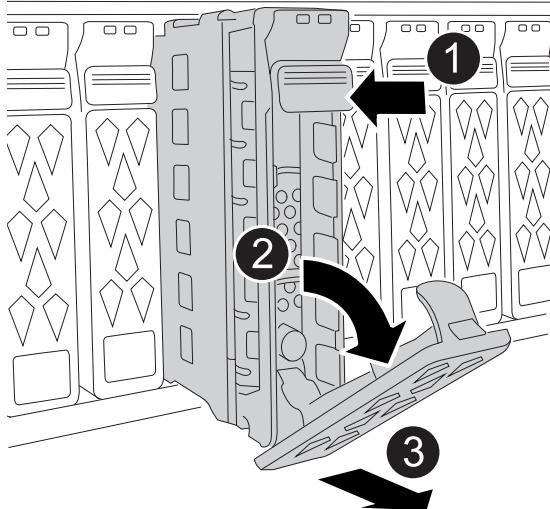
#### Step 2: Remove the drives from the impaired chassis

You need to remove all of the drives and any drive blanks from the impaired chassis so that later in the procedure you can install them in the replacement chassis.

1. Gently remove the bezel from the front of the storage system.
2. Remove the drives and any drive blanks:



Keep track of what drive bay each drive and drive blank was removed from because they must be installed in the same drive bays in the replacement chassis.



1	Press the release button on the drive face to open the cam handle.
2	Rotate the cam handle downward to disengage the drive from the midplane.
3	<p>Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand.</p> <p>When removing a drive, always use two hands to support its weight.</p> <p><span style="border: 1px solid #ccc; border-radius: 50%; padding: 2px 5px; font-size: 14px;">i</span> Because drives are fragile, minimize handling to avoid damaging them.</p>

3. Set the drives aside on a static-free cart or table.

#### Step 2: Replace the chassis from within the equipment rack or system cabinet

You remove the impaired chassis from the equipment rack or system cabinet, install the replacement chassis, install the drives, any drive blanks, and then install the bezel.

1. Remove the screws from the impaired chassis mount points.

Set the screws aside to use later in this procedure.



If the storage system shipped in a NetApp system cabinet, you must remove additional screws at the rear of the chassis before the chassis can be removed.

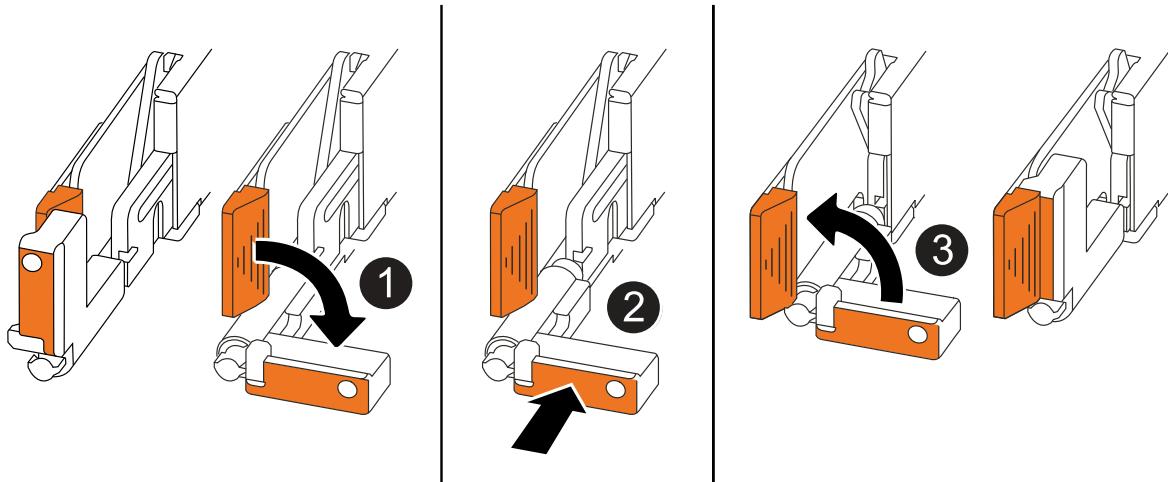
2. Using two people or a power lift, remove the impaired chassis from the equipment rack or system cabinet by sliding it off the rails, and then set it aside.
3. Using two people, install the replacement chassis into the equipment rack or system cabinet by sliding it onto the rails.
4. Secure the front of the replacement chassis to the equipment rack or system cabinet using the screws you removed from the impaired chassis.

#### Step 4: Install the controllers and drives

Install the controllers and drives into the replacement chassis and reboot the controllers.

##### About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when installing a controller, and can be used as a reference for the rest of the controller installation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis and push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

##### 1. Insert one of the controllers into the chassis:

- Align the back of the controller with the opening in the chassis.
- Firmly push on the handles until the controller meets the midplane and is fully seated in the chassis.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

- Rotate the controller handles up and lock in place with the tabs.

##### 2. Recable the controller, as needed, except for the power cords.

##### 3. Repeat these steps to install the second controller into the chassis.

##### 4. Install the drives and any drive blanks you removed from the impaired chassis into the replacement chassis:



The drives and drive blanks must be installed in the same drive bays in the replacement chassis.

- With the cam handle in the open position, use both hands to insert the drive.

- b. Gently push until the drive stops.
- c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

- d. Repeat the process for the remaining drives.

5. Install the bezel.

6. Reconnect the power cords to the power supplies (PSU) in the controllers.

Once power is restored to a PSU, the status LED should be green.



The controllers begin to boot as soon as the power is restored.

If you are reconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"> <li>1. Plug the power cord into the PSU.</li> <li>2. Secure the power cord with the power cord retainer.</li> </ol>
DC PSU	<ol style="list-style-type: none"> <li>1. Plug the D-SUB DC power cord connector into the PSU.</li> <li>2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.</li> </ol>

7. If controllers boot to the LOADER prompt, reboot the controllers:

`boot_ontap`

8. Turn AutoSupport back on:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

### What's next?

After you've replaced the impaired chassis and reinstalled the components into it, you need to [complete the chassis replacement](#).

### Complete chassis replacement - ASA C30

Verify the HA state of the chassis and then return the failed part to NetApp to complete the final step in the ASA C30 chassis replacement procedure.

#### Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your storage system configuration.

1. In Maintenance mode, from either controller, display the HA state of the local controller and chassis:

```
ha-config show
```

The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your storage system configuration:

a. Set the HA state for the chassis:

```
ha-config modify chassis HA-state
```

The value for HA-state should be *ha*.

The value for HA-state can be one of the following:

\* **ha**

\* *mcc* (not supported in ASA)

b. Confirm that the setting has changed:

```
ha-config show
```

3. If you have not already done so, recable the rest of your storage system.

## Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

## Controller

### Controller replacement workflow - ASA C30

Get started with replacing the controller in your ASA C30 storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.

1

#### [Review the controller replacement requirements](#)

To replace the controller, you must meet certain requirements.

2

#### [Shut down the impaired controller](#)

Shut down or take over the impaired controller so that the healthy controller continues to serve data from the impaired controller storage.

3

#### [Replace the controller](#)

Replacing the controller includes removing the impaired controller, moving FRU components to the replacement controller, installing the replacement controller in the chassis, setting the time and date, and then recabling.

4

#### [Restore and verify the system configuration](#)

Verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

5

## Give back the controller

Transfer the ownership of storage resources back to the replacement controller.

6

## Complete controller replacement

Verify the LIFs, check cluster health, and return the failed part to NetApp.

### Requirements to replace the controller - ASA C30

Before replacing the controller in your ASA C30 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.

You must review the requirements and considerations for the controller replacement procedure.

#### Requirements

- All shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- You must replace a controller with a controller of the same model type. You cannot upgrade your system by just replacing the controller.
- You cannot change any drives or shelves as part of this procedure.
- You must always capture the controller's console output to a text log file.

The console output provides you with a record of the procedure you can use to troubleshoot issues you might encounter during the replacement process.

#### Considerations

- It is important that you apply the commands in this procedure to the correct controller:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.

#### What's next?

After you've reviewed the requirements to replace the impaired controller, you need to [shut down the impaired controller](#).

### Shut down the impaired controller - ASA C30

Shut down the impaired controller in your ASA C30 storage system to prevent data loss and ensure system stability when replacing the controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

## About this task

- If you have a SAN system, you must have checked event messages (cluster kernel-service show) for the impaired controller SCSI blade. The cluster kernel-service show command (from priv advanced mode) displays the node name, **quorum status** of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

## Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter *y* when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller:  <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> The <i>-halt true</i> parameter brings you to the LOADER prompt.

## What's next?

After you've shut down the impaired controller, you need to [replace the controller](#).

## Replace the controller - ASA C30

Replace the controller in your ASA C30 storage system when a hardware failure requires

it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

#### About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

#### Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

#### Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

#### Steps

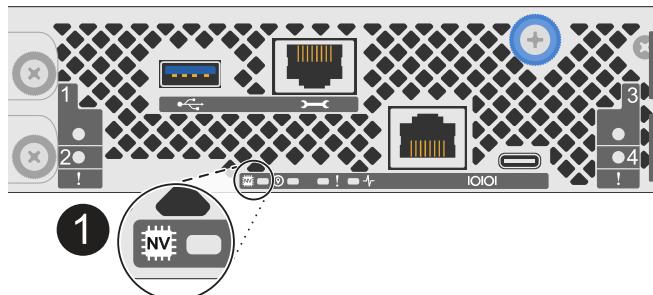
1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.

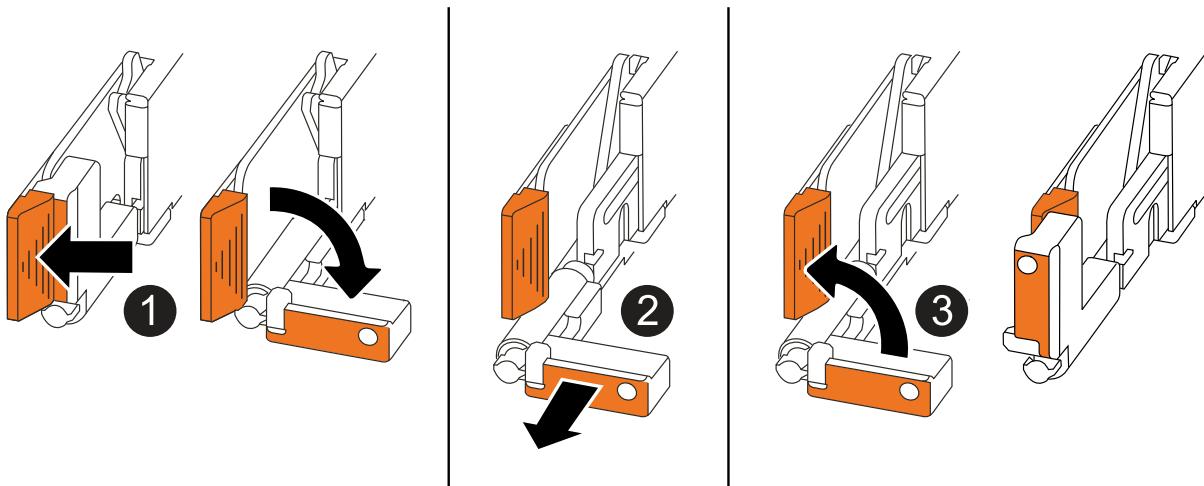
If you are disconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"><li>1. Open the power cord retainer.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>
DC PSU	<ol style="list-style-type: none"><li>1. Unscrew the two thumb screws on the D-SUB DC power cord connector.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	<ul style="list-style-type: none"><li>• Pull the handles towards you to unseat the controller from the midplane.</li></ul> <p>As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling.</p> <ul style="list-style-type: none"><li>• Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.</li></ul>
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

## **Step 2: Move the power supply**

Move the power supply (PSU) to the replacement controller.

### 1. Move the PSU from the impaired controller:

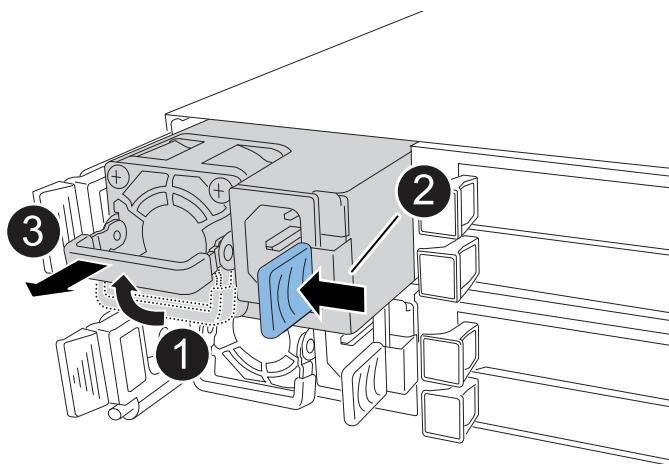
Make sure the left side controller handle is in the upright position to allow you access to the PSU.

## Option 1: Move an AC PSU

To move an AC PSU, complete the following steps.

### Steps

1. Remove the AC PSU from the impaired controller:



1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	Pull the PSU out of the controller while using your other hand to support its weight.   The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.

2. Insert the PSU into the replacement controller:

- a. Using both hands, support and align the edges of the PSU with the opening in the controller.
- b. Gently push the PSU into the controller until the locking tab clicks into place.

A PSU will only properly engage with the internal connector and lock in place one way.



To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

- c. Rotate the handle down, so it is out of the way of normal operations.

## Option 2: Move a DC PSU

To move a DC PSU, complete the following steps.

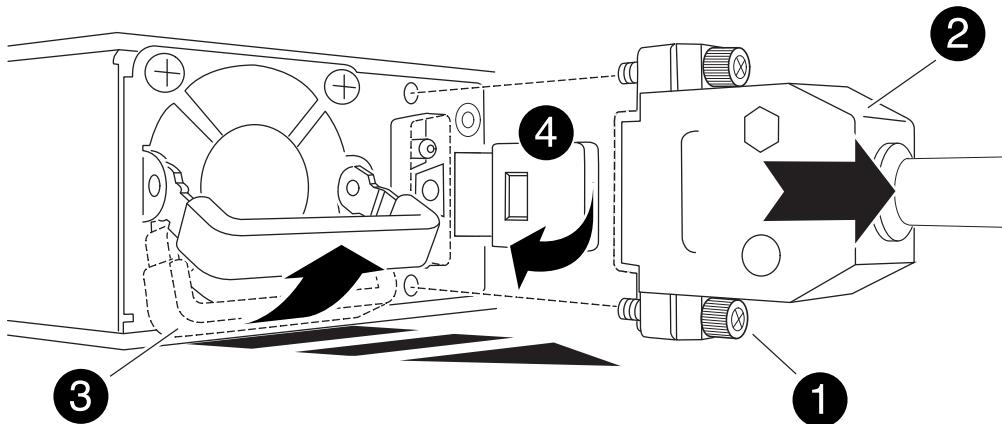
### Steps

1. Remove the DC PSU from the impaired controller:

- a. Rotate the handle up, to its horizontal position, and then grasp it.
- b. With your thumb, press the terracotta tab to release the locking mechanism.
- c. Pull the PSU out of the controller while using your other hand to support its weight.



The PSU is short. Always use two hands to support it when removing it from the controller so that it does not swing free from the controller and injure you.



1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

## 2. Insert the PSU into the replacement controller:

- a. Using both hands, support and align the edges of the PSU with the opening in the controller.
- b. Gently slide the PSU into the controller until the locking tab clicks into place.

A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.



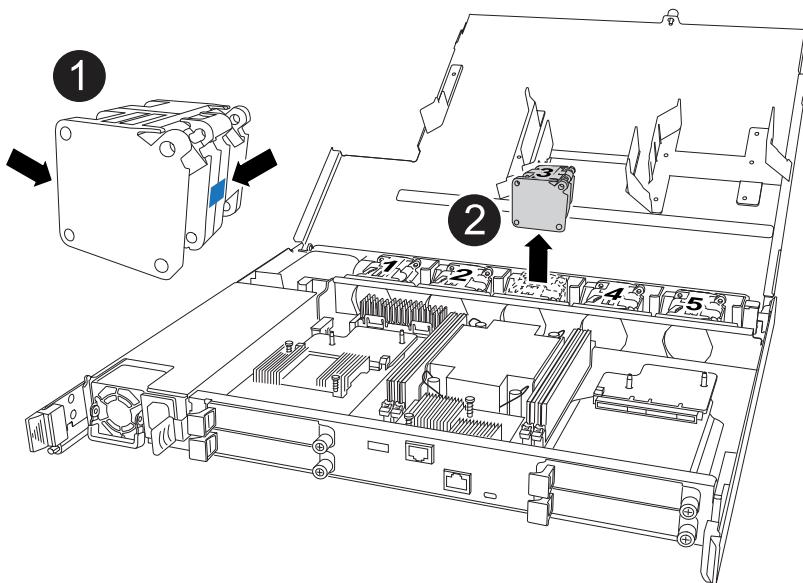
To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

- c. Rotate the handle down, so it is out of the way of normal operations.

## Step 3: Move the fans

Move the fans to the replacement controller.

### 1. Remove one of the fans from the impaired controller:



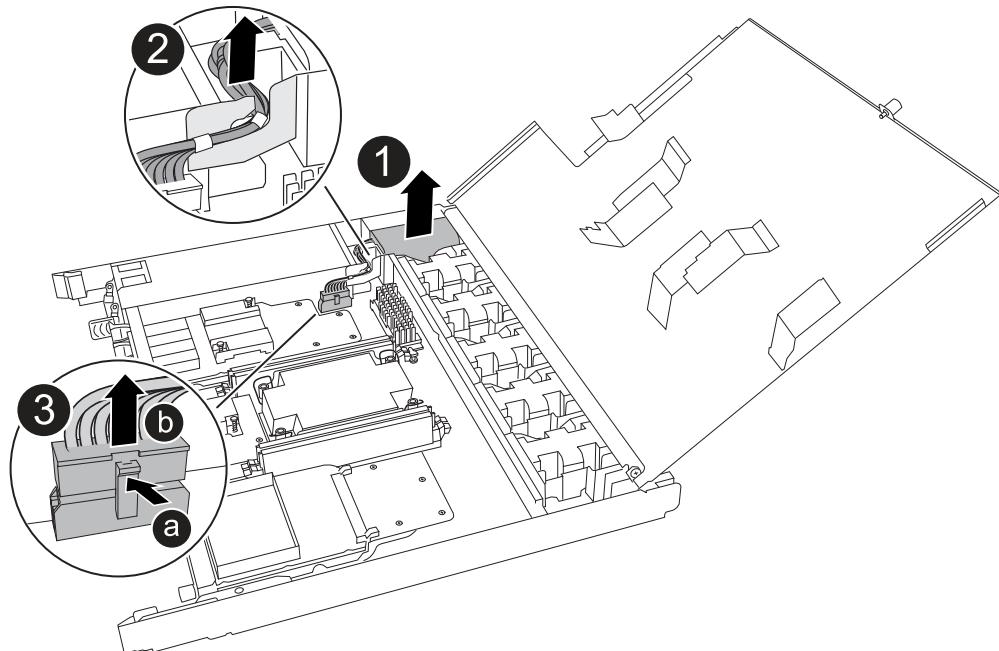
1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

2. Insert the fan into the replacement controller by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.
3. Repeat these steps for the remaining fans.

#### Step 4: Move the NV battery

Move the NV battery to the replacement controller.

1. Remove the NV battery from the impaired controller:



1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	<ol style="list-style-type: none"> <li>1. Push in and hold the tab on the connector.</li> <li>2. Pull the connector up and out of the socket.</li> </ol> <p>As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.</p>

2. Install the NV battery into the replacement controller:

- a. Plug the wiring connector into its socket.
- b. Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
- c. Place the NV battery into the compartment.

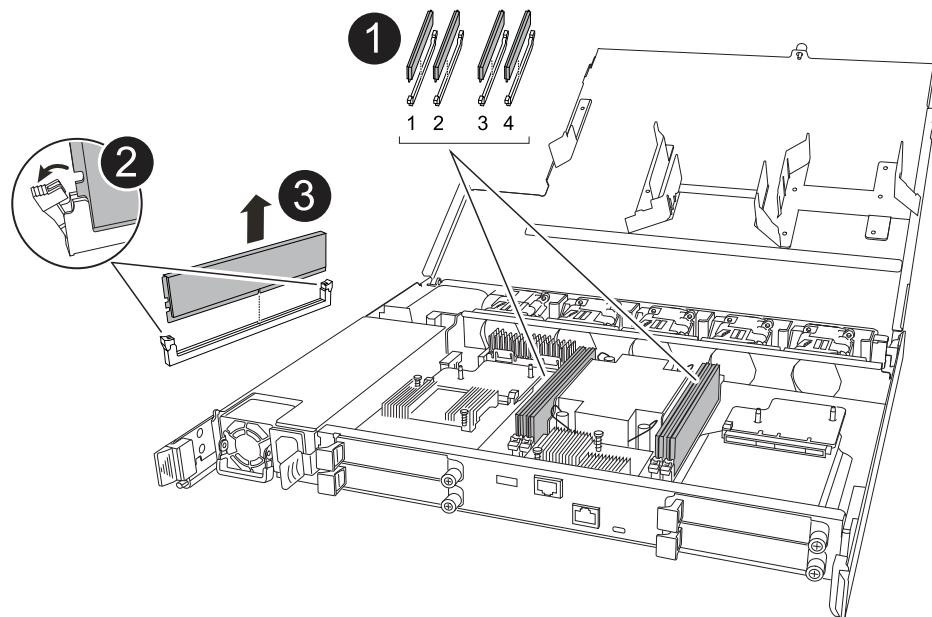
The NV battery should sit flush in its compartment.

#### Step 5: Move system DIMMs

Move the DIMMs to the replacement controller.

If you have DIMM blanks, you do not need to move them, the replacement controller should come with them installed.

1. Remove one of the DIMMs from the impaired controller:



1	<p>DIMM slot numbering and positions.</p> <p> Depending on your storage system model, you will have two or four DIMMs.</p>
2	<ul style="list-style-type: none"> <li>• Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller in the proper orientation.</li> <li>• Eject the DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot.</li> </ul> <p> Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.</p>
3	<p>Lift the DIMM up and out of the slot.</p> <p>The ejector tabs remain in the open position.</p>

2. Install the DIMM in the replacement controller:

- Make sure that the DIMM ejector tabs on the connector are in the open position.
- Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

When inserted correctly, the DIMM goes in easily but fits tightly in the slot. If not, reinsert the DIMM.

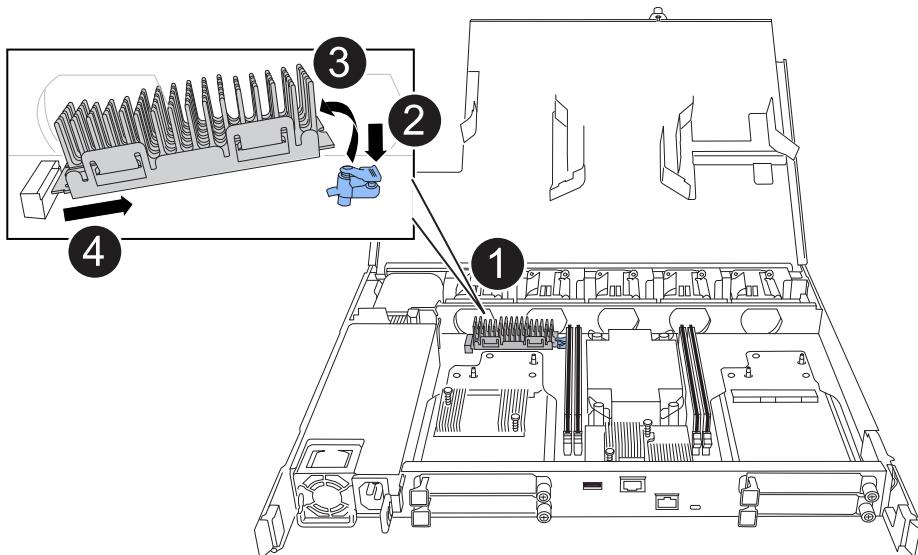
- Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
- Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.

3. Repeat these steps for the remaining DIMMs.

**Step 6: Move the boot media**

Move the boot media to the replacement controller.

1. Remove the boot media from the impaired controller:



1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

2. Install the boot media into the replacement controller:

- Slide the socket end of the boot media into its socket.
- At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

**Step 7: Move the I/O modules**

Move the I/O modules and any I/O blanking modules to the replacement controller.

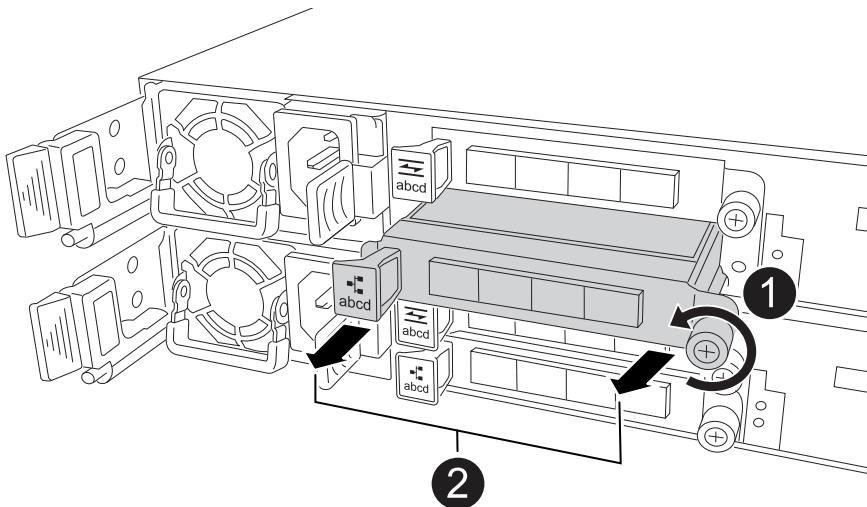
1. Unplug cabling from one of the I/O modules.

Make sure to label the cables so that you know where they came from.

2. Remove the I/O module from the impaired controller:

Make sure that you keep track of which slot the I/O module was in.

If you are removing the I/O module in slot 4, make sure the right side controller handle is in the upright position to allow you access to the I/O module.



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

3. Install the I/O module into the replacement controller:

a. Align the I/O module with the edges of the slot.

b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O Module.

c. Turn the thumbscrew clockwise to tighten.

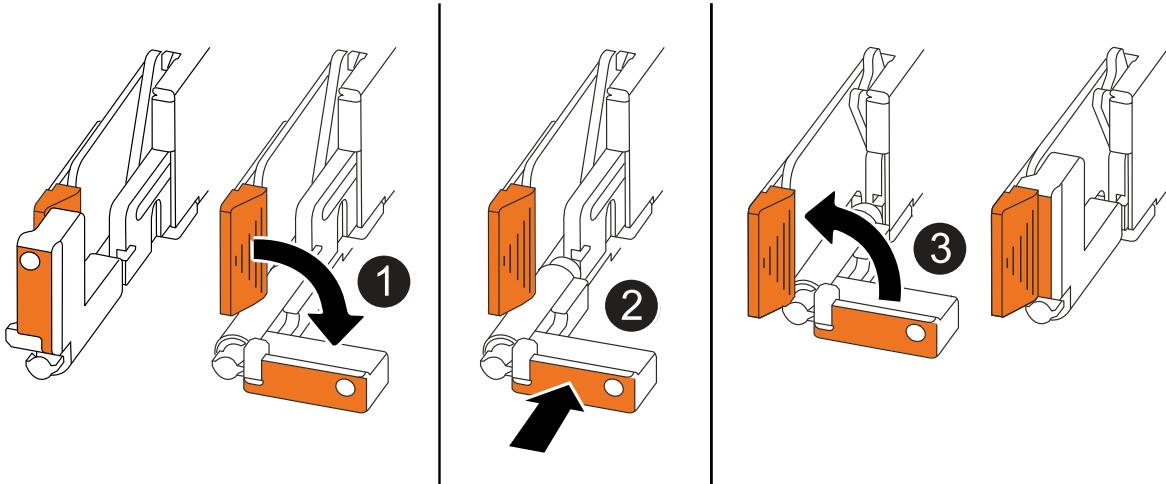
4. Repeat these steps to move the remaining I/O modules and any I/O blanking modules to the replacement controller.

**Step 8: Install the controller**

Reinstall the controller into the chassis and reboot it.

**About this task**

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

## Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.

2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.



Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:

a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

b. Rotate the controller handles up and lock in place with the tabs.



The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Take the controller to the LOADER prompt by pressing CTRL-C to abort AUTOBOOT.
6. Set the time and date on the controller:

Make sure you are at the controller's LOADER prompt.

- a. Display the date and time on the controller:

```
show date
```



Time and date default is in GMT. You have the option to display in local time and in 24hr mode.

- b. Set the current time in GMT:

```
set time hh:mm:ss
```

You can get the current GMT from the healthy node:

```
date -u
```

- c. Set the current date in GMT:

```
set date mm/dd/yyyy
```

You can get the current GMT from the healthy node:

```
date -u
```

7. Recable the controller as needed.

8. Reconnect the power cord to the power supply (PSU).

Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"><li>1. Plug the power cord into the PSU.</li><li>2. Secure the power cord with the power cord retainer.</li></ol>
DC PSU	<ol style="list-style-type: none"><li>1. Plug the D-SUB DC power cord connector into the PSU.</li><li>2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.</li></ol>

#### What's next?

After you've replaced the impaired controller, you need to [restore the system configuration](#).

#### Restore and verify the system configuration - ASA C30

Verify that the controller's HA configuration is active and functioning correctly in your ASA C30 storage system, and confirm that the system's adapters list all the paths to the disks.

## Step 1: Verify HA config settings

You must verify the HA state of the controller and, if necessary, update the state to match your storage system configuration.

1. Boot to maintenance mode:

```
boot_ontap maint
```

- a. Enter `y` when you see *Continue with boot?*.

If you see the *System ID mismatch* warning message, enter `y`.

2. Enter `sysconfig -v` and capture the display contents.



If you see *PERSONALITY MISMATCH* contact customer support.

3. From the `sysconfig -v` output, compare the adapter card information with the cards and locations in the replacement controller.

4. Verify that all components display the same HA state:

```
ha-config show
```

The HA state should be the same for all components.

5. If the displayed system state of the controller does not match your storage system configuration, set the HA state for the controller:

```
ha-config modify controller ha
```

The value for the HA state can be one of the following:

- ha
- mcc (not supported)
- mccip (not supported in ASA systems)
- non-ha (not supported)

6. Confirm that the setting has changed:

```
ha-config show
```

## Step 2: Verify disk list

1. Verify that the adapter lists the paths to all disks:

```
storage show disk -p
```

If you see any issues, check cabling and reseat cables.

2. Exit Maintenance mode:

halt

#### What's next?

After you've restored and verified your system configuration, you need to [give back the controller](#).

#### Give back the controller - ASA C30

Return control of storage resources to the replacement controller so your ASA C30 storage system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption, Onboard Key Manager (OKM) encryption, or External Key Manager (EKM) encryption.

## No encryption

Return the impaired controller to normal operation by giving back its storage.

### Steps

1. From the LOADER prompt, enter `boot_ontap`.
2. Press `<enter>` when console messages stop.
  - If you see the *login* prompt, go to the next step at the end of this section.
  - If you see *Waiting for giveback*, press the `<enter>` key, log into the partner node, and then go to the next step at the end of this section.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

## Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

### Steps

1. From the LOADER prompt, enter `boot_ontap maint`.
2. Boot to the ONTAP menu from the LOADER prompt `boot_ontap` menu and select option 10.
3. Enter the OKM passphrase.



You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option 1 for normal boot.
6. Press `<enter>` when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as admin.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`
9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
  - a. Move the console cable back to the replacement controller.
  - b. Synchronize missing keys: `security key-manager onboard sync`



If you encounter errors, contact [NetApp Support](#).



You are prompted for the cluster-wide passphrase of OKM for the cluster.

c. Verify status of the keys: `security key-manager key query -restored false`

The output should show no results when properly synchronized.

If the output shows results (the key IDs of keys that are not present in the system's internal key table), contact [NetApp Support](#).

11. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
12. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
13. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

### External key manager (EKM)

Reset encryption and return the controller to normal operation.

#### Steps

1. If the root volume is encrypted with External Key Manager and the console cable is connected to the replacement node, enter `boot_ontap` menu and select option 11.
2. If these questions appear, answer y or n as appropriate:

Do you have a copy of the `/cfcard/kmip/certs/client.crt` file? {y/n}

Do you have a copy of the `/cfcard/kmip/certs/client.key` file? {y/n}

Do you have a copy of the `/cfcard/kmip/certs/CA.pem` file? {y/n}

Do you have a copy of the `/cfcard/kmip/servers.cfg` file? {y/n}

Do you know the KMIP server address? {y/n}

Do you know the KMIP port? {y/n}



Contact [NetApp Support](#) if you have issues.

3. Supply the information for:
  - The client certificate (`client.crt`) file contents
  - The client key (`client.key`) file contents
  - The KMIP server CA(s) (`CA.pem`) file contents
  - The IP address for the KMIP server
  - The port for the KMIP server
4. Once the system processes, you see the Boot Menu. Select '1' for normal boot.
5. Check the takeover status: `storage failover show`
6. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
7. If automatic giveback was disabled, reenable it: `storage failover modify -node local`

```
-auto-giveback true
```

8. If AutoSupport is enabled, restore/unsuppress automatic case creation: system node autosupport invoke -node \* -type all -message MAINT=END

## What's next?

After you've transferred the ownership of storage resources to the replacement controller, you need to [complete the controller replacement](#) procedure.

## Complete controller replacement - ASA C30

To complete the controller replacement for your ASA C30 storage system, first restore the NetApp Storage Encryption configuration (if necessary) and install the required licenses on the new controller. Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, register the new controller's serial number and then return the failed part to NetApp.

### Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

#### Before you begin

If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in [Post Motherboard Replacement Process to update Licensing on ONTAP platforms](#). If you are unsure of the initial ONTAP release for your system, see [NetApp Hardware Universe](#) for more information.

#### About this task

- Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

- The license keys must be in the 28-character format.
- You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.
- If the node is in a MetroCluster configuration and all nodes at a site have been replaced, license keys must be installed on the *replacement* node or nodes prior to switchback.

#### Steps

1. If you need new license keys, obtain replacement license keys on the [NetApp Support site](#) in the My Support section under Software licenses.



The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key license-key...`
3. Remove the old licenses, if desired:
  - a. Check for unused licenses: `license clean-up -unused -simulate`
  - b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

### Step 2: Verify LIFs, register the serial number, and check cluster health

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

#### Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`  
If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`
2. Register the system serial number with NetApp Support.
  - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
  - If AutoSupport is not enabled, call [NetApp Support](#) to register the serial number.
3. Check the health of your cluster. See the [How to perform a cluster health check with a script in ONTAP](#) KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

### Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

## Replace a DIMM - ASA C30

Replace a DIMM in your ASA C30 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.

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

#### Before you begin

- Ensure all other components in the storage system are working correctly; if not, contact [NetApp Support](#) before continuing."

- You must replace the failed FRU component with a replacement FRU component you received from your provider.

## About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

## Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

## About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, `quorum status` of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

## Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.

If the impaired controller is displaying...	Then...
Waiting for giveback...	Press Ctrl-C, and then respond <b>y</b> when prompted.
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 <b>-halt true</b> parameter brings you to the LOADER prompt.

## Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

### Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

### Steps

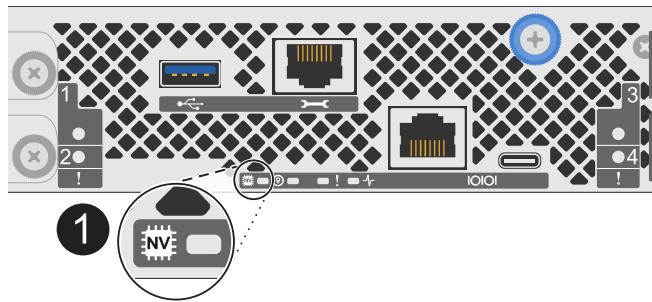
1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.

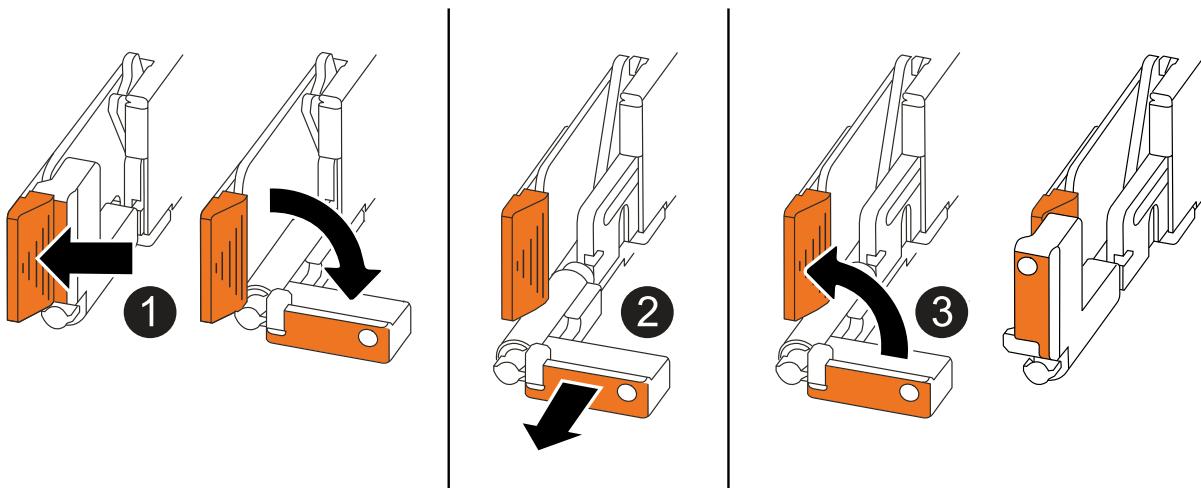
If you are disconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"> <li>1. Open the power cord retainer.</li> <li>2. Unplug the power cord from the PSU and set it aside.</li> </ol>
DC PSU	<ol style="list-style-type: none"> <li>1. Unscrew the two thumb screws on the D-SUB DC power cord connector.</li> <li>2. Unplug the power cord from the PSU and set it aside.</li> </ol>

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	<ul style="list-style-type: none"> <li>• Pull the handles towards you to unseat the controller from the midplane.</li> </ul> <p>As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling.</p> <ul style="list-style-type: none"> <li>• Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.</li> </ul>
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

### Step 3: Replace a DIMM

To replace a DIMM, locate the faulty DIMM inside the controller and follow the specific sequence of steps.

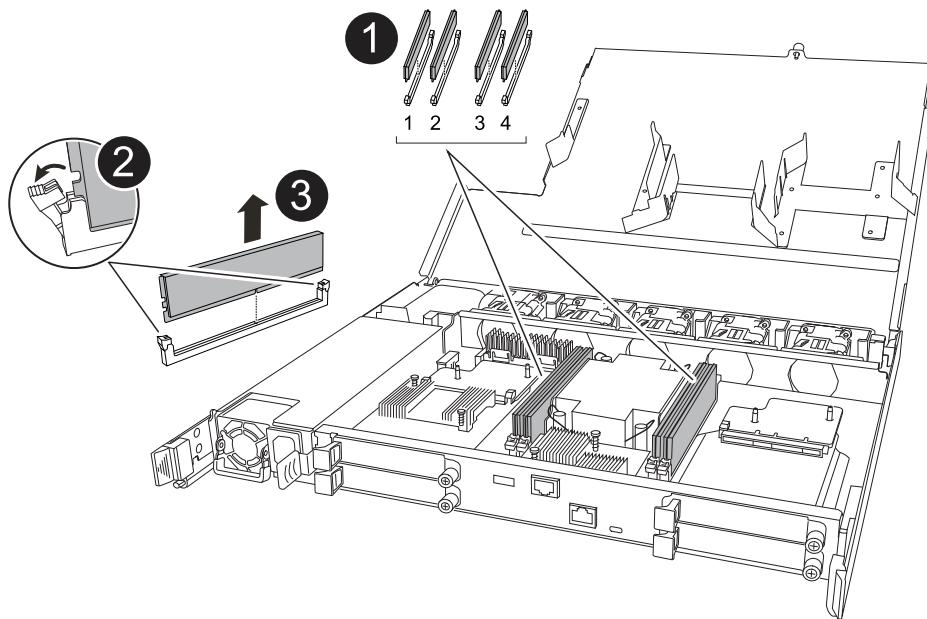
#### Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller and identify the faulty DIMM.



Consult either the [Netapp Hardware Universe](#) or the FRU map on the cover of the controller for exact DIMM locations.

3. Remove the faulty DIMM:



<b>1</b>	<p>DIMM slot numbering and positions.</p> <p> Depending on your storage system model you will have two or four DIMMs.</p>
<b>2</b>	<ul style="list-style-type: none"><li>• Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM using the same orientation.</li><li>• Eject the faulty DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot.</li></ul> <p> Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.</p>
<b>3</b>	<p>Lift the DIMM up and out of the slot.</p> <p>The ejector tabs remain in the open position.</p>

#### 4. Install the replacement DIMM:

- Remove the replacement DIMM from its antistatic shipping bag.
- Make sure that the DIMM ejector tabs on the connector are in the open position.
- Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

When inserted correctly, the DIMM goes in easily but fits tightly in the slot. Reinsert the DIMM if you feel it is not inserted correctly.

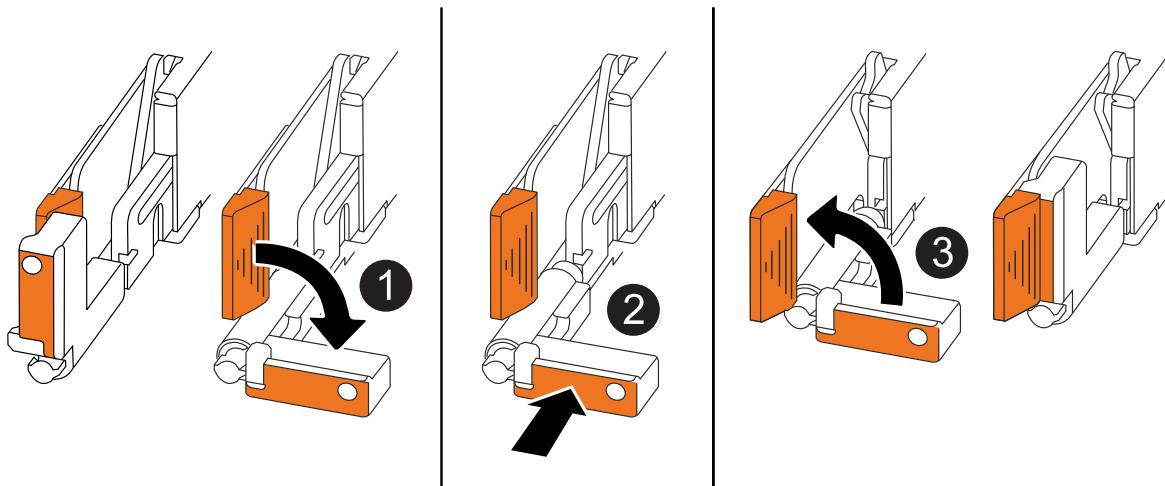
- Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
- Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.

### Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

#### About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

#### Steps

- Close the controller cover and turn the thumbscrew clockwise until tightened.
- Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.



Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:

- a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

- b. Rotate the controller handles up and lock in place with the tabs.



The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.

6. Reconnect the power cord to the power supply (PSU).

Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"><li>1. Plug the power cord into the PSU.</li><li>2. Secure the power cord with the power cord retainer.</li></ol>
DC PSU	<ol style="list-style-type: none"><li>1. Plug the D-SUB DC power cord connector into the PSU.</li><li>2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.</li></ol>

7. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

8. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

## Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

## Replace a drive - ASA C30

Replace a drive in your ASA C30 storage system when a drive fails or requires an upgrade. The replacement process involves identifying the faulty drive, safely removing it, and installing a new drive to ensure continued data access and system performance.

You can replace a failed drive nondisruptively while I/O is in progress.

### Before you begin

- The drive that you are installing must be supported by your storage system.

#### [NetApp Hardware Universe](#)

- If self-encrypting drive (SED) authentication is enabled, you must use the SED replacement instructions in the ONTAP documentation.

Instructions in the ONTAP documentation describe additional steps you must perform before and after replacing an SED.

#### [NetApp encryption overview with the CLI](#)

- All other components in the storage system must be functioning properly, if not, you must contact [NetApp Support](#) before continuing with this procedure.
- Verify that the drive you are removing is failed.

You can verify that the drive is failed by running the `storage disk show -broken` command. The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.



Depending on the drive type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

### About this task

- When replacing a failed drive, you must wait 70 seconds between the removal of the drive and the insertion of the replacement drive to allow the storage system to recognize that a drive was removed.
- The best practice is to have the current version of the Disk Qualification Package (DQP) installed before hot-swapping a drive.

Having the current version of the DQP installed allows your system to recognize and use newly qualified drives. This avoids system event messages about having noncurrent drive information and prevention of drive partitioning because drives are not recognized. The DQP also notifies you of noncurrent drive firmware.

#### [NetApp Downloads: Disk Qualification Package](#)

- The best practice is to have current versions of NVMe shelf module (NSM) firmware and drive firmware on your system before replacing FRU components.



Do not revert firmware to a version that does not support your shelf and its components.

- Drive firmware is automatically updated (nondisruptively) on new drives that have non-current firmware versions.



Drive firmware checks occur every two minutes.

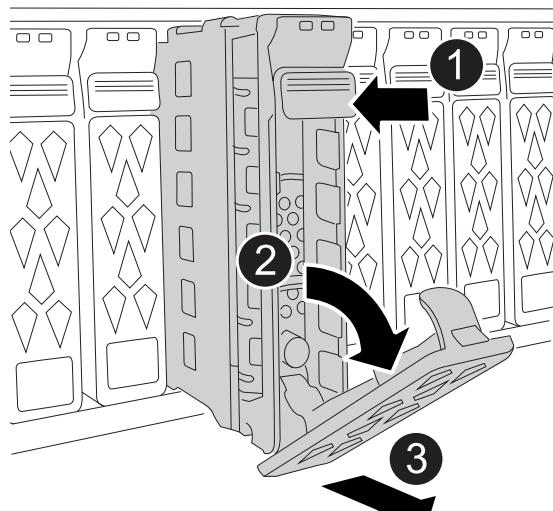
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

## Steps

- Properly ground yourself.
- Remove the bezel from the front of the storage system.
- Physically identify the failed drive.
  - When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the shelf operator display panel and the failed drive illuminate.
  - The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
- Remove the failed drive:



1

Press the release button on the drive face to open the cam handle.

2	Rotate the cam handle downward to disengage the drive from the midplane.
3	<p>Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand.</p> <p>When removing a drive, always use two hands to support its weight.</p> <p>Because drives are fragile, minimize handling to avoid damaging them.</p>

5. Wait a minimum of 70 seconds before inserting the replacement drive.

This allows the system to recognize that a drive was removed.

6. Insert the replacement drive:

- a. With the cam handle in the open position, use both hands to insert the replacement drive.
- b. Gently push until the drive stops.
- c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

7. Verify that the drive's activity (green) LED is illuminated.

When the drive's activity LED is solid, it means that the drive has power. When the drive's activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat Steps 3 through 7.

9. Reinstall the bezel on the front of the storage system.

10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Contact technical support [NetApp Support](#) if you need the RMA number or additional help with the replacement procedure.

## Replace a fan module - ASA C30

Replace a fan module in your ASA C30 storage system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

### About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs

are on or off, you can check their state by entering the `system location-led` show command.

## Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

### About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, [quorum status](#) of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

### Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter *y* when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

## Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

### Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

### Steps

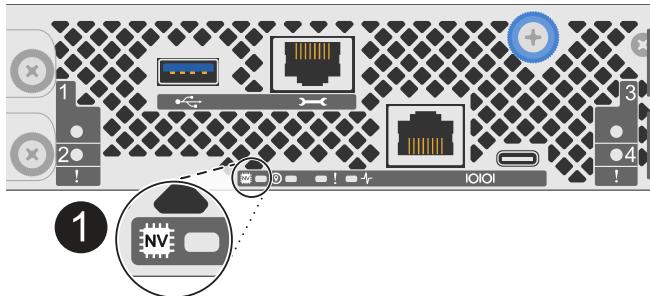
1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.

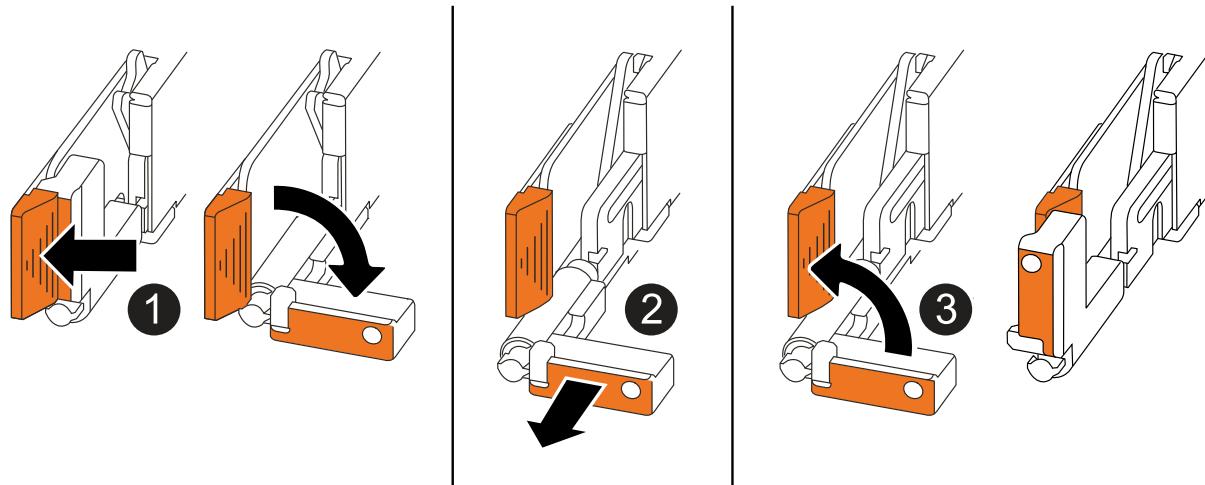
If you are disconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"><li>1. Open the power cord retainer.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>
DC PSU	<ol style="list-style-type: none"><li>1. Unscrew the two thumb screws on the D-SUB DC power cord connector.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	<ul style="list-style-type: none"><li>Pull the handles towards you to unseat the controller from the midplane.</li><li>As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling.</li><li>Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.</li></ul>
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

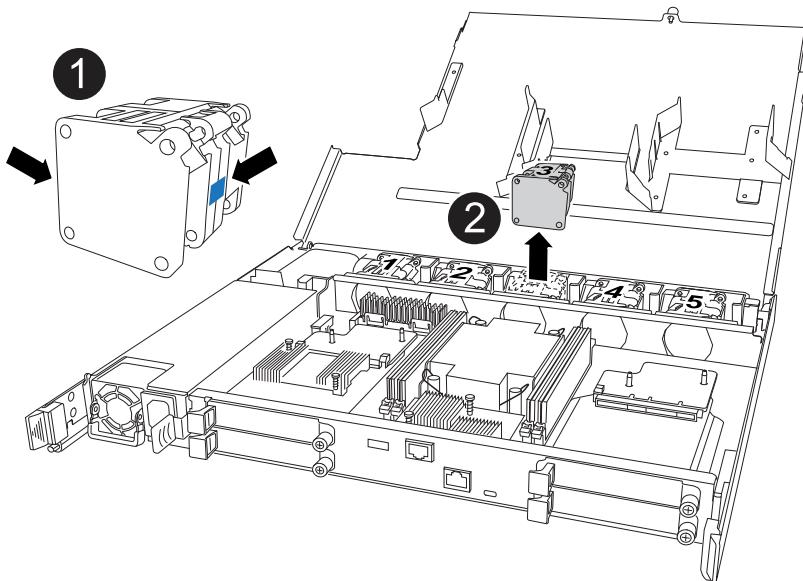
6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

### Step 3: Replace fan

To replace a fan, remove the failed fan and replace it with a new fan.

#### Steps

1. Identify the fan that you must replace by checking the console error messages.
2. Remove the failed fan:



1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

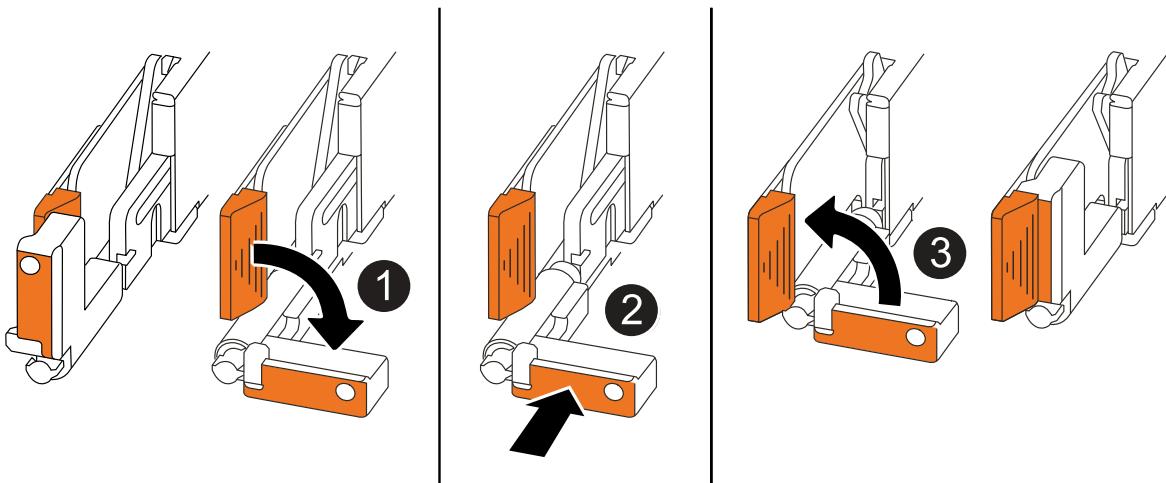
3. Insert the replacement fan by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.

#### Step 4: Reinstall the controller module

Reinstall the controller into the chassis and reboot it.

##### About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

## Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.



Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:

- a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

- b. Rotate the controller handles up and lock in place with the tabs.



The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.

6. Reconnect the power cord to the power supply (PSU).

Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"> <li>1. Plug the power cord into the PSU.</li> <li>2. Secure the power cord with the power cord retainer.</li> </ol>
DC PSU	<ol style="list-style-type: none"> <li>1. Plug the D-SUB DC power cord connector into the PSU.</li> <li>2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.</li> </ol>

7. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

8. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

## Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

## I/O module

### Overview of I/O module maintenance - ASA C30

The ASA C30 storage system offers flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding, hot-swapping, or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your storage system with the same type of I/O module, or with a different type of I/O module. You can hot-swap a cluster and HA I/O module when your storage system meets specific requirements. You can also add an I/O module to a storage system with available slots.

- [Add an I/O module](#)

Adding additional I/O modules can improve redundancy, helping to ensure that the storage system remains operational even if one I/O module fails.

- [Hot-swap an I/O module](#)

You can hot-swap certain I/O modules for an equivalent I/O module to restore the storage system to its optimal operating state. Hot-swapping is done without having to perform a manual takeover.

To use this procedure, your storage system must be running ONTAP 9.17.1 or later and meet specific system requirements.

- [Replace an I/O module](#)

Replacing a failing I/O module can restore the storage system to its optimal operating state.

### Add an I/O module - ASA C30

Add an I/O module to your ASA C30 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 ASA C30 storage system if there are available slots. If all slots are fully populated, you can replace an existing module to add a new one.

## About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

### Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

## About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, `quorum status` of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

## Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.

If the impaired controller is displaying...	Then...
Waiting for giveback...	Press Ctrl-C, and then respond <b>y</b> when prompted.
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 <b>-halt true</b> parameter brings you to the LOADER prompt.</p>

## Step 2: Add the new I/O module

If the storage system has available slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

### Before you begin

- Check the [NetApp Hardware Universe](#) to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in [NetApp Hardware Universe](#) and use the best one available for your I/O module.
- All other components in the storage system must be functioning properly; if not, contact [NetApp Support](#) before continuing with this procedure.

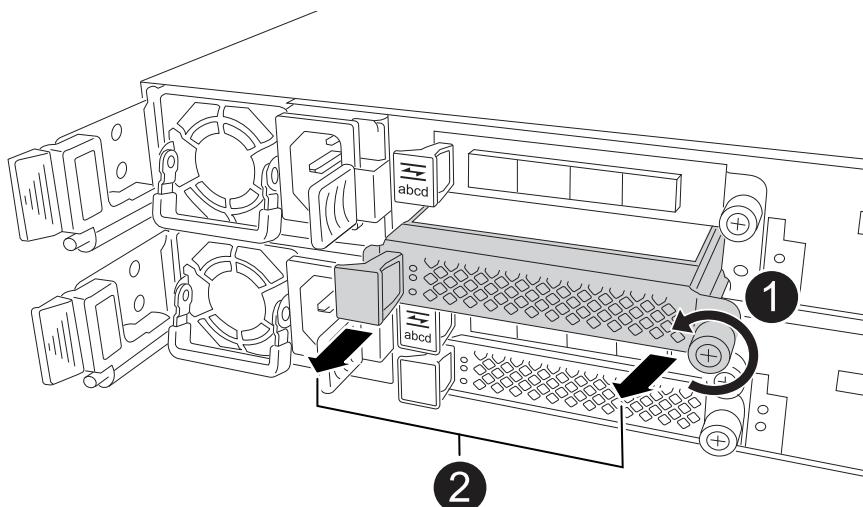
## Add I/O module to an available slot

You can add a new I/O module into a storage system with available slots.

### Steps

1. If you are not already grounded, properly ground yourself.
2. On the impaired controller, remove the I/O blanking module from the target slot.

Unused I/O slots should have blanking module installed to prevent possible thermal issues and for EMC compliance.



1	On the I/O blanking module, turn the thumbscrew counterclockwise to loosen.
2	Pull the I/O blanking module out of the controller using the tab on the left and the thumbscrew.

3. Install the new I/O module:

- a. Align the I/O module with the edges of the controller slot opening.
- b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O module.

- c. Turn the thumbscrew clockwise to tighten.

4. Cable the I/O module to the designated devices.

If you installed a storage I/O module, install and cable your NS224 shelves, as described in [Hot-add workflow](#).

5. Reboot the impaired controller from the LOADER prompt: `bye`

Rebooting the impaired controller also reinitializes the I/O modules and other components.

6. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name.
```

7. Repeat these steps to add an I/O module to the other controller.
8. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation: +

```
system node autosupport invoke -node * -type all -message MAINT=END
```

### Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

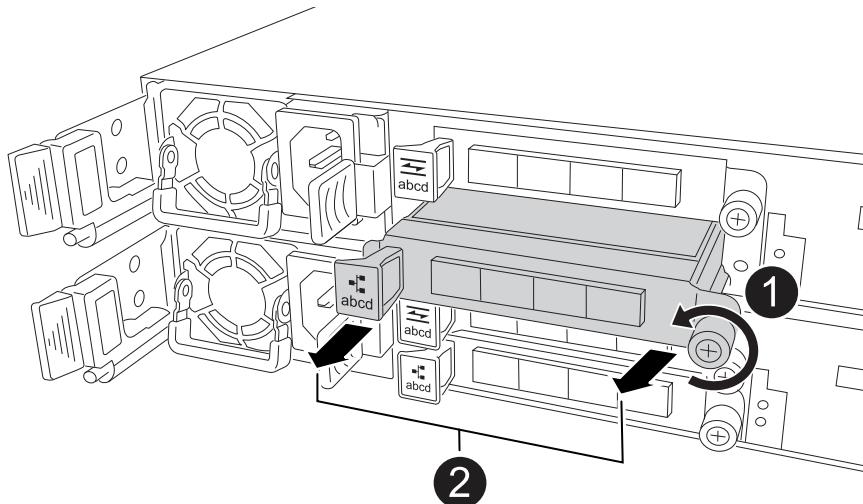
#### About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See <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. On the impaired controller, unplug any cabling on the target I/O module.
3. Remove the target I/O module from the controller:



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the new I/O module into the target slot:

- a. Align the I/O module with the edges of the slot.
- b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O Module.

- c. Turn the thumbscrew clockwise to tighten.

5. Cable the I/O module to the designated devices.

If you installed a storage I/O module, install and cable your NS224 shelves, as described in [Hot-add workflow](#).

6. Repeat the I/O module remove and install steps to add any additional I/O modules in the controller.

7. Reboot the impaired controller from the LOADER prompt:

bye

Rebooting the impaired controller also reinitializes the I/O modules and other components.

8. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

9. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

10. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

11. If you installed a NIC module, specify the usage mode for each port as *network*:

```
storage port modify -node node_name -port port_name -mode network
```

12. Repeat these steps for the other controller.

## Hot swap an I/O module - ASA C30

You can hot swap an Ethernet I/O module in your ASA C30 storage system if a module fails and your storage system meets all ONTAP version requirements.

To hot swap an I/O module, ensure your storage system meets the ONTAP version requirements, 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

- Hot-swapping the I/O module means that you do not have to perform a manual takeover before replacing the failed I/O module.
- Apply commands to the correct controller and I/O slot when you are hot-swapping the I/O module:
  - The *impaired controller* is the controller on which you are hot-swapping 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.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

## Step 1: Ensure the storage system meets the procedure requirements

To use this procedure, your storage system must be running ONTAP 9.17.1 or later, and your storage system must meet all requirements for the version of ONTAP your storage system is running.



If your storage system is not running ONTAP 9.17.1 or later, or does not meet all requirements for the version of ONTAP your storage system is running, you cannot use this procedure, you must use the [replace an I/O module procedure](#).

## ONTAP 9.17.1 or 9.18.1RC

- You are hot-swapping a failed cluster and HA I/O module in slot 4 with an equivalent I/O module. You cannot change the I/O module type.
- The controller with the failed cluster and HA I/O module (the impaired controller) must have already taken over the healthy partner controller. The takeover should have occurred automatically if the I/O module has failed.

For two-node clusters, the storage system cannot discern which controller has the failed I/O module, so either controller might initiate the takeover. Hot swapping is only supported when the controller with the failed I/O module (the impaired controller) has taken over the healthy controller. Hot-swapping the I/O module is the only way to recover without an outage.

You can verify that the impaired controller successfully took over the healthy controller by entering the `storage failover show` command.

If you are not sure which controller has the failed I/O module, contact [NetApp Support](#).

- Your storage system configuration must have only one cluster and HA I/O module located in slot 4, not two cluster and HA I/O modules.
- Your storage system must be a two-node (switchless or switched) cluster configuration.
- All other components in the storage system must be functioning properly; if not, contact [NetApp Support](#) before continuing with this procedure.

## ONTAP 9.18.1GA or later

- 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 controllers 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. Unplug the cables from the failed I/O module.

Make sure to label the cables so you can reconnect them to the same ports later in this procedure.

 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. As needed for the version of ONTAP your storage system is running and the state of the controllers, disable automatic giveback:

ONTAP version	If...	Then...
9.17.1 or 9.18.1RC	If the impaired controller took over the healthy controller automatically	<p>Disable automatic giveback:</p> <ol style="list-style-type: none"><li>1. Enter the following command from the console of the impaired controller</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>
9.18.1GA or later	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>

ONTAP version	If...	Then...
9.18.1GA or later	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:

- Enter the following command:

```
system controller slot module remove -node impaired_node_name -slot
slot_number
```

- Enter *y* when you see the prompt *Do you want to continue?*

For example, the following command prepares the failed module in slot 4 on node 2 (the impaired controller) for removal, and displays a message that it is safe to remove:

```
node2::> system controller slot module remove -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered off for removal.

Do you want to continue? {y|n}: y

The module has been successfully removed from service and powered
off. It can now be safely removed.
```

6. Verify the failed I/O module is powered off:

```
system controller slot module show
```

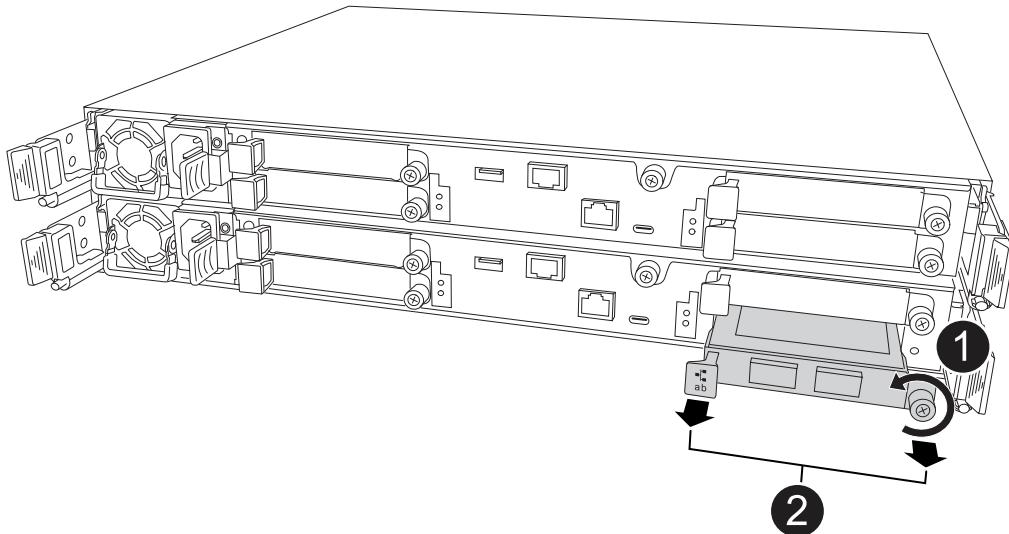
The output should show *powered-off* in the *status* column for the failed module and its slot number.

### Step 3: Hot swap the failed I/O module

Hot swap the failed I/O module with an equivalent I/O module:

#### Steps

- If you are not already grounded, properly ground yourself.
- Remove the failed I/O module from the impaired controller:



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew on the right.

3. Install the replacement I/O module:

- Align the I/O module with the edges of the slot.
- Gently push the I/O module all the way into the slot, making sure to properly seat the I/O module into the connector.

You can use the tab on the left and the thumbscrew on the right to push in the I/O module.

- Turn the thumbscrew clockwise to tighten.

4. Cable the replacement I/O module.

**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:

- Enter the following command:

```
system controller slot module insert -node impaired_node_name -slot
slot_number
```

- 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 4 on node 2 (the impaired controller) online, and displays a message that the process was successful:

```
node2::> system controller slot module insert -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered on and initialized.

Do you want to continue? {y|n}: `y`

The module has been successfully powered on, initialized and placed
into service.
```

2. Verify that each port on the 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 indicating each port on the I/O module initiated successfully.

For example, the following output shows initialization succeeded for I/O ports e4b and e4a:

```
node2::> event log show -event *hotplug.init*

Time           Node           Severity      Event
-----
-----
7/11/2025 16:04:06  node2      NOTICE       hotplug.init.success:
Initialization of ports "e4b" in slot 4 succeeded

7/11/2025 16:04:06  node2      NOTICE       hotplug.init.success:
Initialization of ports "e4a" in slot 4 succeeded

2 entries were displayed.
```

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 a I/O module in slot 4:

```
node2::> system controller config show -node local -slot 4

Node: node2
Sub- Device/
Slot slot Information
-----
4      - Dual 40G/100G Ethernet Controller CX6-DX
          e4a MAC Address: d0:39:ea:59:69:74 (auto-100g_cr4-fd-
up)
          QSFP Vendor:      CISCO-BIZLINK
          QSFP Part Number: L45593-D218-D10
          QSFP Serial Number: LCC2807GJFM-B
          e4b MAC Address: d0:39:ea:59:69:75 (auto-100g_cr4-fd-
up)
          QSFP Vendor:      CISCO-BIZLINK
          QSFP Part Number: L45593-D218-D10
          QSFP Serial Number: LCC2809G26F-A
          Device Type:      CX6-DX PSID(NAP0000000027)
          Firmware Version: 22.44.1700
          Part Number:      111-05341
          Hardware Revision: 20
          Serial Number:    032403001370
```

#### Step 5: Restore the storage system to normal operation

Restore your storage system to normal operation by giving back storage to the 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:

ONTAP version	If...	Then...
9.17.1 or 9.18.1RC	If the impaired controller took over the healthy controller automatically	<ol style="list-style-type: none"> <li>1. Return the healthy controller to normal operation by giving back its storage:  <code>storage failover giveback -ofnode <i>healthy_node_name</i></code></li> <li>2. Restore automatic giveback from the console of the impaired controller:  <code>storage failover modify -node local -auto-giveback true</code></li> </ol>
9.18.1GA or later	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 <i>controller that was taken over_name</i></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>
9.18.1GA or later	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 server and ports: `network interface show -is-home false`

If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

3. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

#### Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

#### Replace an I/O module - ASA C30

Replace an I/O module in your ASA C30 storage system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

Use this procedure to replace a failed I/O module.

## Before you begin

All other components in the storage system must be functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

## About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

### Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

## About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, [quorum status](#) of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

## Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

## Step 2: Replace a failed I/O module

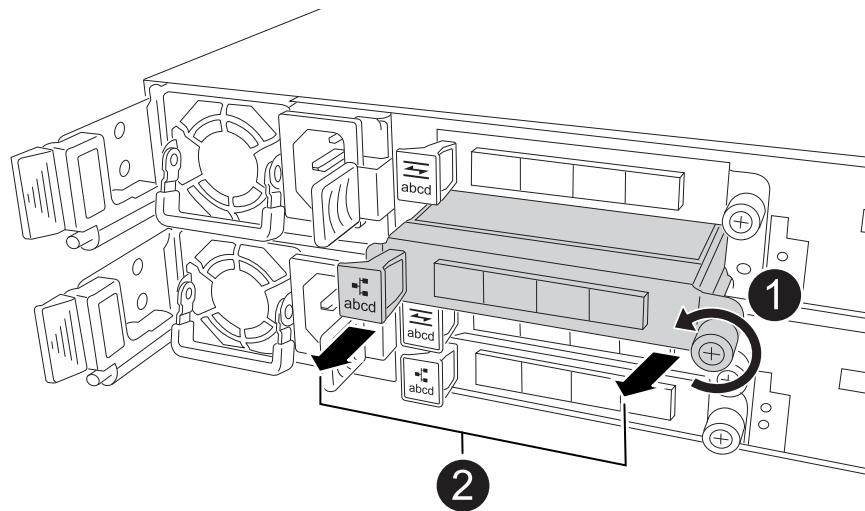
To replace a failed I/O module, locate it in the controller and follow the specific sequence of steps.

### Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug cabling from the failed I/O module.

Make sure to label the cables so that you know where they came from.

3. Remove the failed I/O module from the controller:



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the replacement I/O module into the target slot:
  - a. Align the I/O module with the edges of the slot.

- b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O Module.

- c. Turn the thumbscrew clockwise to tighten.

## 5. Cable the I/O module.

### Step 3: Reboot the controller

After you replace an I/O module, you must reboot the controller.

#### Steps

1. Reboot the controller from the LOADER prompt: `bye`

Rebooting the impaired controller also reinitializes the I/O modules and other components.

2. Return the node to normal operation: `storage failover giveback -ofnode impaired_node_name`
3. Restore automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback true`

### Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

## Replace the NV battery - ASA C30

Replace the NV battery in your ASA C30 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.

To replace the NV battery, you must remove the controller, remove the faulty battery, install the replacement battery, and then reinstall the controller.

#### Before you begin

All other components in the storage system must be functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

#### About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

## Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

### About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, `quorum status` of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

### Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <code>y</code> when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller:  <code>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</code>  The <code>-halt true</code> parameter brings you to the LOADER prompt.

## Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

### Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

### Steps

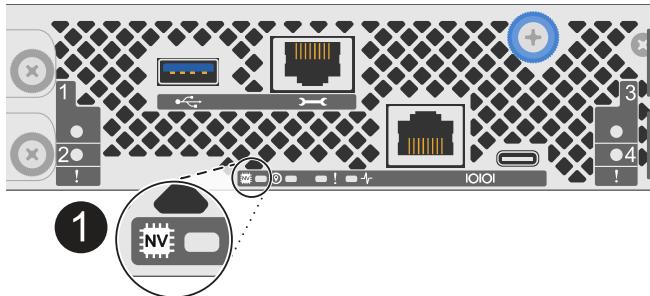
1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.

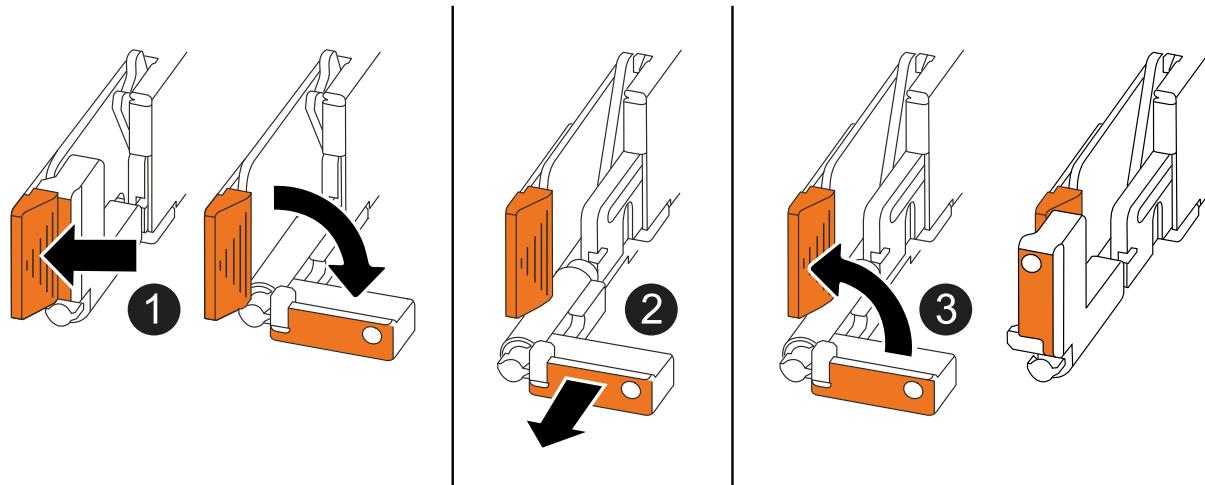
If you are disconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"><li>1. Open the power cord retainer.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>
DC PSU	<ol style="list-style-type: none"><li>1. Unscrew the two thumb screws on the D-SUB DC power cord connector.</li><li>2. Unplug the power cord from the PSU and set it aside.</li></ol>

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	<ul style="list-style-type: none"><li>Pull the handles towards you to unseat the controller from the midplane.</li><li>As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling.</li><li>Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.</li></ul>
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

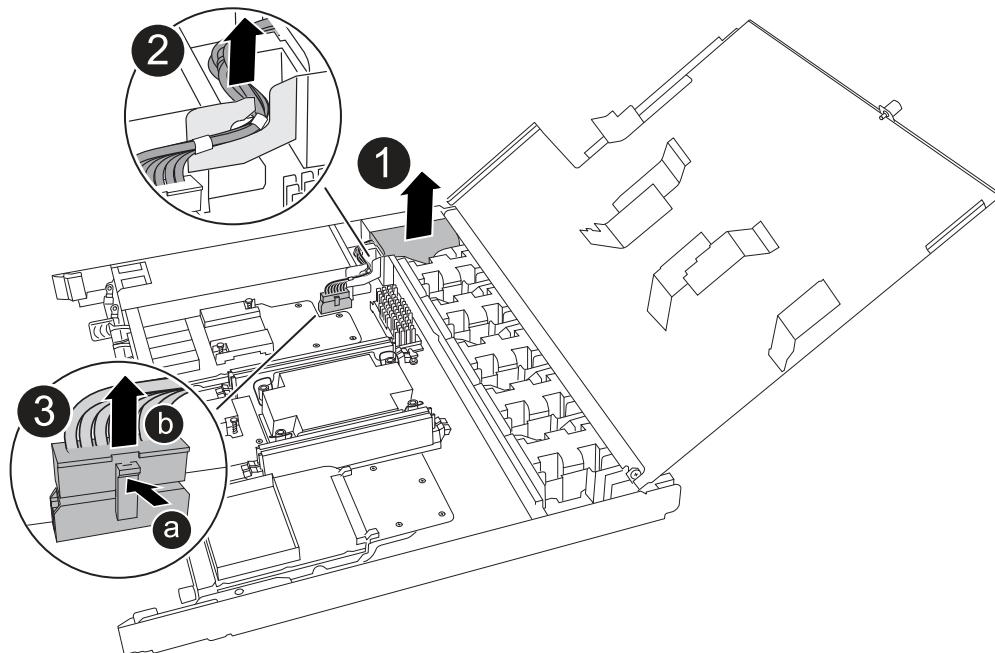
6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

### Step 3: Replace the NV battery

Remove the failed NV battery from the controller and install the replacement NV battery.

#### Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the NV battery.
3. Remove the NV battery:



1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	<ol style="list-style-type: none"> <li>Push in and hold the tab on the connector.</li> <li>Pull the connector up and out of the socket.</li> </ol> <p>As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.</p>

4. Install the replacement NV battery:

- Remove the replacement battery from its package.
- Plug the wiring connector into its socket.
- Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
- Place the NV battery into its compartment.

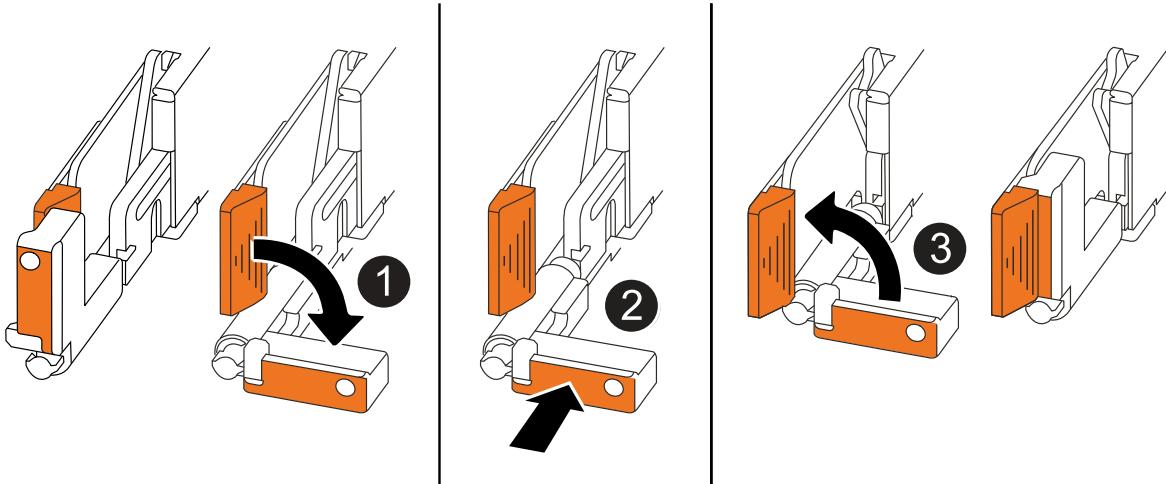
The NV battery should sit flush in its compartment.

#### Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

#### About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

## Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.

2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.



Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:

a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

b. Rotate the controller handles up and lock in place with the tabs.



The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"> <li>1. Plug the power cord into the PSU.</li> <li>2. Secure the power cord with the power cord retainer.</li> </ol>
DC PSU	<ol style="list-style-type: none"> <li>1. Plug the D-SUB DC power cord connector into the PSU.</li> <li>2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.</li> </ol>

7. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

8. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

## Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

## Replace a power supply - ASA C30

Replace an AC or DC power supply unit (PSU) in your ASA C30 storage system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cord, replacing the faulty PSU, and then reconnecting it to the power source.

### About this task

- This procedure is written for replacing one PSU at a time.

The PSUs are redundant and hot-swappable. You do not have to shut down the controller to replace a PSU.

- **IMPORTANT:** Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.
- Use the appropriate procedure for your type of PSU: AC or DC.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected

storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

## Option 1: Hot-swap an AC PSU

To replace an AC PSU, complete the following steps.

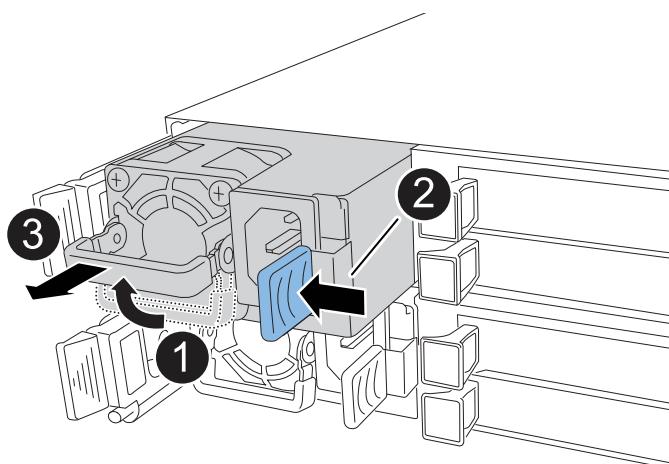
### Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the power cord from the PSU by opening the power cord retainer, and then unplug the power cord from the PSU.



PSUs do not have a power switch.

4. Remove the PSU:



1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	<p>Pull the PSU out of the controller while using your other hand to support its weight.</p> <p> The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.</p>

5. Install the replacement PSU:

- a. Using both hands, support and align the edges of the PSU with the opening in the controller.
- b. Gently push the PSU into the controller until the locking tab clicks into place.

A PSU will only properly engage with the internal connector and lock in place one way.



To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

- c. Rotate the handle down, so it is out of the way of normal operations.
6. Reconnect the power cord to the PSU and secure the power cord with the power cord retainer.

Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

### Option 2: Hot-swap a DC PSU

To replace a DC PSU, complete the following steps.

#### Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the PSU:

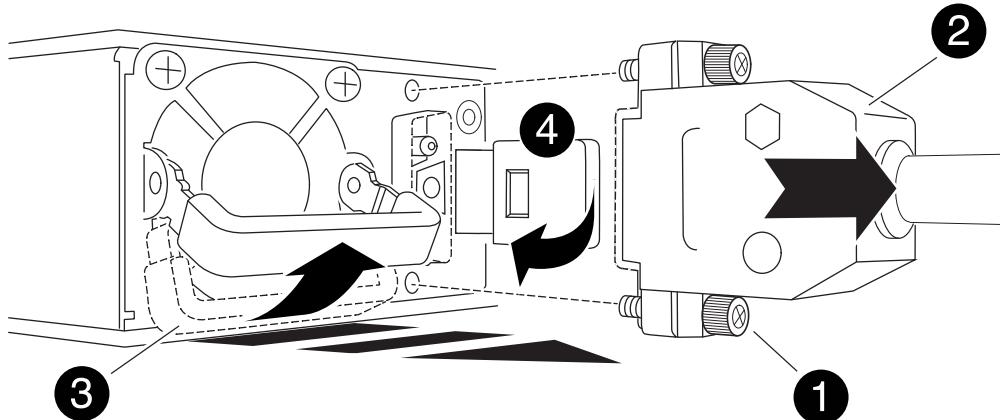


PSUs do not have a power switch.

- a. Unscrew the two thumb screws on the D-SUB DC power cord connector.
- The illustration and table in step 4 shows the two thumb screws (item #1) and the D-SUB DC power cord connector (item #2).
- b. Unplug the cord from the PSU and set it aside.
4. Remove the PSU:
  - a. Rotate the handle up, to its horizontal position, and then grasp it.
  - b. With your thumb, press the terracotta tab to release the locking mechanism.
  - c. Pull the PSU out of the controller while using your other hand to support its weight.



The PSU is short. Always use two hands to support it when removing it from the controller so that it does not swing free from the controller and injure you.



1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

5. Insert the replacement PSU:

- Using both hands, support and align the edges of the PSU with the opening in the controller.
- Gently slide the PSU into the controller until the locking tab clicks into place.

A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.



To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

- Rotate the handle down, so it is out of the way of normal operations.

6. Reconnect the D-SUB DC power cord:

Once power is restored to the PSU, the status LED should be green.

- Plug the D-SUB DC power cord connector into the PSU.
- Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

## Replace the real-time clock battery - ASA C30

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your ASA C30 storage system to ensure that services and applications relying on accurate time synchronization remain operational.

You replace the real-time clock (RTC) battery in the controller so that your storage system's services and applications that depend on accurate time synchronization continue to function.

### Before you begin

All other components in the storage system must be functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

### About this task

- You can use this procedure with all versions of ONTAP supported by your storage system.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

### Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

#### About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, **quorum status** of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

#### Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <code>y</code> when prompted.

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

## Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

### Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

### Steps

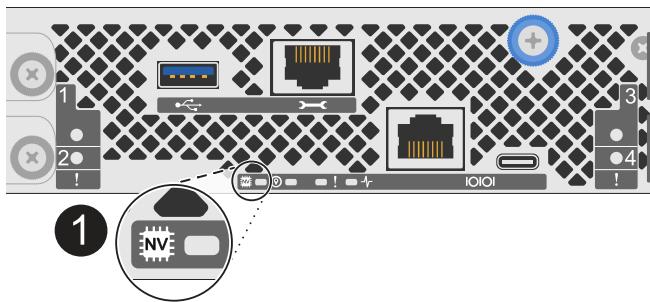
1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.

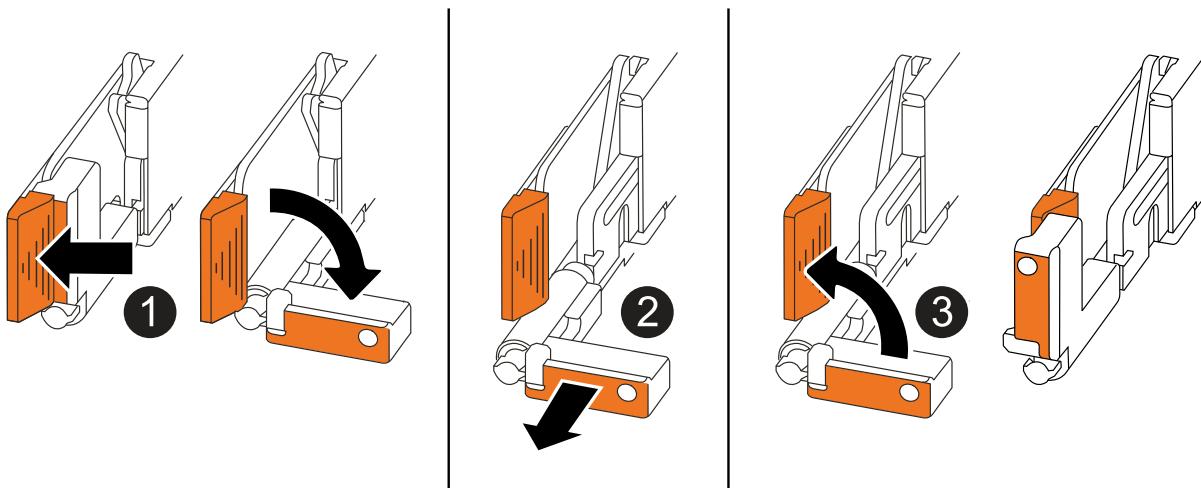
If you are disconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"> <li>1. Open the power cord retainer.</li> <li>2. Unplug the power cord from the PSU and set it aside.</li> </ol>
DC PSU	<ol style="list-style-type: none"> <li>1. Unscrew the two thumb screws on the D-SUB DC power cord connector.</li> <li>2. Unplug the power cord from the PSU and set it aside.</li> </ol>

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	<ul style="list-style-type: none"> <li>• Pull the handles towards you to unseat the controller from the midplane.</li> </ul> <p>As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling.</p> <ul style="list-style-type: none"> <li>• Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.</li> </ul>
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

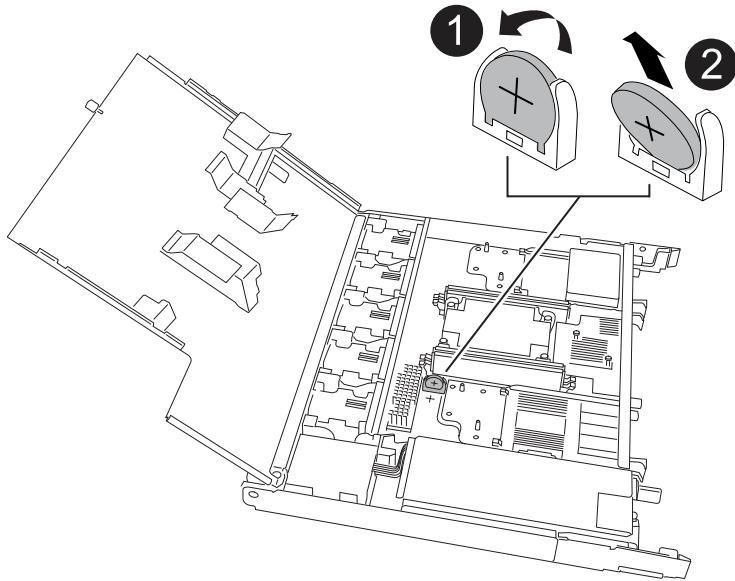
6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

### Step 3: Replace the RTC battery

Remove the failed RTC battery and install the replacement RTC battery.

#### Steps

1. Locate the RTC battery.
2. Remove the RTC battery:



1	Gently rotate the RTC battery at an angle away from its holder.
2	Lift the RTC battery out of its holder.

3. Install the replacement RTC battery:

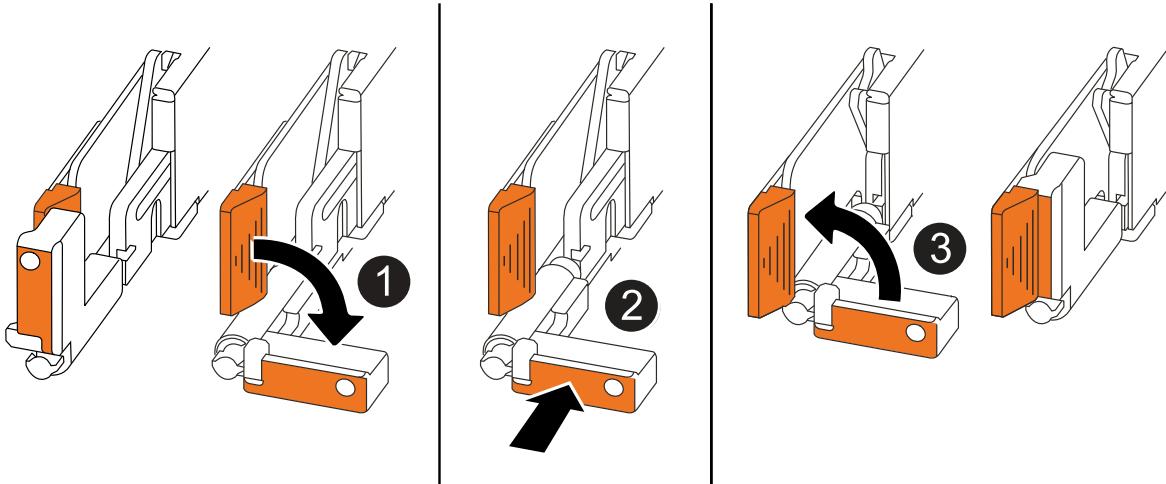
- a. Remove the replacement battery from the antistatic shipping bag.
- b. Position the battery so that the plus sign on the battery faces out to correspond with the plus sign on the motherboard.
- c. Insert the battery into the holder at an angle, and then push it into an upright position so it is fully seated in the holder.
- d. Visually inspect the battery to make sure that it is completely seated in its holder and that the polarity is correct.

### Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

#### About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

## Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.

2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.



Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:

a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

b. Rotate the controller handles up and lock in place with the tabs.



The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a...	Then...
AC PSU	<ol style="list-style-type: none"> <li>1. Plug the power cord into the PSU.</li> <li>2. Secure the power cord with the power cord retainer.</li> </ol>
DC PSU	<ol style="list-style-type: none"> <li>1. Plug the D-SUB DC power cord connector into the PSU.</li> <li>2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.</li> </ol>

7. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

8. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

## Step 5: Reset the time and date on the controller

After replacing the RTC battery, inserting the controller and powering on first BIOS reset, you will see the following error messages:



RTC date/time error. Reset date/time to default

RTC power failure error

These messages are expected and you can continue with this procedure.

1. On the healthy controller, check the date and time with the `cluster date show` command.
 

If your storage system stops at the boot menu, select the option for `Reboot node` and respond `y` when prompted, then boot to LOADER by pressing `Ctrl-C`.
2. On the impaired controller, at the LOADER prompt, check the time and date: `cluster date show`
  - a. If necessary, modify the date: `set date mm/dd/yyyy`
  - b. If necessary, set the time, in GMT: `set time hh:mm:ss`
  - c. Confirm the date and time.
3. At the LOADER prompt, enter `bye` to reinitialize the I/O modules, other components, and let the controller reboot.
4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`

5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

#### **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.

## 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.