



AFF A70 and AFF A90 systems

Install and maintain

NetApp
July 01, 2024

Table of Contents

- AFF A70 and AFF A90 systems 1
 - Install and setup 1
 - Maintain 18

AFF A70 and AFF A90 systems

Install and setup

Install setup workflow - AFF A70, AFF A90

Follow these workflow steps to install, set up, and deploy your AFF A70, AFF A90 storage system.

1

Prepare for install

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

2

Install hardware

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

3

Cable hardware

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

4

Power on storage system

To power on your system, first power on each NS224 shelf and assign a unique shelf ID to ensure each shelf is distinct within the setup. Then, power on the controllers. Finally, use the provided securing device to secure the power cables on each power supply.

5

Complete storage system setup

To complete system setup, access the web-based interface, System Manager, by pointing a browser to the controller's IP address. A setup wizard helps you get started with cluster configuration. As an alternative, you can also use the command line interface (CLI).

Prepare for install - AFF A70, AFF A90

Prepare to install your platform.

Before you begin


- Use the [cluster setup worksheet](#) to record the values that you need when you [configure your cluster](#). If a default value is provided, you can use that value or else enter your own.
- Confirm you have all of the additional equipment and tools that you need to complete the hardware installation.

- Phillips #2 screwdriver
 - Flashlight
 - ESD strap
 - A paperclip or narrow tipped ball point pen for setting NS224 shelf IDs
 - Laptop or console with a USB/serial connection
 - Access to a web browser to configure your storage system
- You should be aware of the safety concerns associated with the weight of the platform.

A fully loaded A70-90 platform can weigh up to 151.68 lbs (68.8 kg) and requires four people to lift or use a hydraulic lift.

- You should be aware of the safety concerns associated with the weight of the shelf.

A fully loaded shelf can weigh up to 66.78 lbs (30.29 kg) and requires two people to lift or use a hydraulic lift. Avoid removing shelf components (from the front or rear of the shelf) to reduce the shelf weight, because shelf weight will become unbalanced.



- Airflow through the storage system enters from the front where the bezel or end caps are installed and exhausts out the rear where the ports are located.

Step 1: Prepare the site

Make sure that the site and the cabinet or rack you plan to use meet the specifications for your storage system.

Steps

1. Confirm that your site meets the environmental and electrical requirements for your platform as shown in the [NetApp Hardware Universe](#).
2. Make sure you have adequate rack space:
 - 4U for the platform chassis
 - 2U for each NS224 shelf
3. Install any required network switches near the platform.

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

Step 2: Unpack the boxes

Unpack all boxes and compare the contents to the items on the packing slip.

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



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

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

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

Hardware

- Platform with onboard drives installed
- Shelf (optional)
- Rail kits with instructions (optional)
- Two Cable management devices
- Bezel

Cables

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

Step 3: Register your hardware

Register your platform to get support benefits.

Steps

1. Locate the system serial number for your platform.

You can find the number on the packing slip, in your confirmation email, or on the controller after you unpack it.



2. Go to the NetApp Support Site at mysupport.netapp.com.
3. Register the hardware:

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

Install the hardware - AFF A70, AFF A90

The following instructions describe how to install the hardware in an empty cabinet or telco rack. Skip this step if your cabinet is pre-populated.

Before you begin


- You should review the [Safety Notices](#) document, and understand the precautions for moving and installing hardware.
- You should have the instructions packaged with the rail kit.
- You should be aware of the safety concerns associated with the weight of the platform and shelf.

Steps

1. Install the rail kits for the platform and shelves, as needed, using the instructions included with the kits.
2. Install and secure your platform in the cabinet or telco rack:
 - a. Position the platform onto the rails in the middle of the cabinet or telco rack, and then support the platform from the bottom and slide it into place.
 - b. Secure the platform to the cabinet or telco rack using the included mounting screws.
3. Install the shelf or shelves:
 - a. Position the back of the shelf onto the rails, and then support the shelf from the bottom and slide it into the cabinet or telco rack.

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

- b. Secure the shelf to the cabinet or telco rack using the included mounting screws.
4. Attach the two cable management devices to the rear of the controllers.




5. Attach the bezel to the front of the platform.

Cable the hardware - AFF A70, AFF A90

Cable your controllers and shelves according to your specific storage system configuration.

Before you begin

Check the illustration arrow in the cabling diagrams for the proper cable connector pull-tab orientation.



- As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn the cable head over and try again.
- If connecting to an optical switch, insert the SFP into the controller port before cabling to the port.

Step 1: Cable controllers to your network

Cable the controllers to your network as either a two-node switchless cluster or a switched cluster.

Before you begin

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

About this task

These procedures show common configurations. Keep in mind that the specific cabling depends on the components ordered for your storage system. For comprehensive configuration and slot priority details, see [NetApp Hardware Universe](#).

Option 1: Cable switchless cluster

Cable your switchless cluster connections.

Steps


1. Cable the Cluster/HA interconnect ports e1a to e1a and e7a to e7a with the Cluster/HA interconnect cable.



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



Cluster/HA interconnect cables



2. Cable the Ethernet module ports to your host network.

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


- a. Cable ports e9a and e9b to your Ethernet data network switch as shown.



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


100 GbE cable





b. Cable your 10/25 GbE host network switches.


4-ports, 10/25 GbE Host



3. Cable the controller management (wrench) ports to the management network switches with 1000BASE-T RJ-45 cables.



1000BASE-T RJ-45 cables



DO NOT plug in the power cords yet.

Option 2: Cable switched cluster

Cable your switched cluster connections.

Steps


1. Make the following cabling connections:



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

- a. Cable port e1a on Controller A and port e1a on Controller B to cluster network switch A.
- b. Cable port e7a on Controller A and port e7a on Controller B to cluster network switch B.

100 GbE cable



2. Cable the Ethernet module ports to your host network.


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

- a. Cable ports e9a and e9b to your Ethernet data network switch as shown.




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

100 GbE cable



b. Cable your 10/25 GbE host network switches.

4-ports, 10/25 GbE Host




3. Cable the controller management (wrench) ports to the management network switches with 1000BASE-T RJ-45 cables.



1000BASE-T RJ-45 cables

Controller A

Controller B



To management network switches



DO NOT plug in the power cords yet.

Step 2: Cable controllers to shelves

Cable your controllers to the shelf or shelves.


These procedures show how to cable your controllers to one shelf and to two shelves. You can connect your controllers to a maximum of four shelves.

Option 1: Cable to one NS224 shelf


Cable each controller to the NSM modules on the NS224 shelf. The graphics show cabling from each of the controllers: Controller A cabling in blue and Controller B cabling in yellow.

Steps

1. Connect controller A port e11a to NSM A port e0a.
2. Connect controller A port e11b to port NSM B port e0b.



3. Connect controller B port e11a to NSM B port e0a.
4. Connect controller B port e11b to NSM A port e0b.




Option 2: Cable to two NS224 shelves


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

Steps

1. On controller A, cable the following connections:
 - a. Connect port e11a to shelf 1, NSM A port e0a.
 - b. Connect port e11b to shelf 2, NSM B port e0b.
 - c. Connect port e8a to shelf 2, NSM A port e0a.
 - d. Connect port e8b to shelf 1, NSM B port e0b.



2. On controller B, cable the following connections:
 - a. Connect port e11a to shelf 1, NSM B port e0a.
 - b. Connect port e11b to shelf 2, NSM A port e0b.
 - c. Connect port e8a to shelf 2, NSM B port e0a.
 - d. Connect port e8b to shelf 1, NSM A port e0b.



Power on the storage system - AFF A70, AFF A90

After connecting the network and any shelf cables, you are ready to apply power to the controllers and any additional shelves.



External NS224 shelves are optional.

Step 1: Power on the shelf and assign shelf ID

In ONTAP storage systems, each NS224 shelf is distinguished by a unique shelf ID. This ID ensures that the shelf is distinct within your storage system setup. By default, shelf IDs are typically assigned as '00' and '01', but you may need to adjust these IDs to maintain uniqueness across your storage system.



The controller chassis shelf ID is hard-coded to '00'. Do not use shelf ID '00' for any NS224 shelves attached to the platform.

About this task


- A valid shelf ID is 00 through 99.
- You must power cycle a shelf (unplug both power cords, wait the appropriate amount of time, and then plug them back in) for the shelf ID to take effect.

Steps

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

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

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



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

3. Change the first number of the shelf ID:

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

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



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

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

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

The first number continues to blink.

4. Change the second number of the shelf ID:

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

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

The first number on the digital display stops blinking.

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

The second number continues to blink.

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

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

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

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

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

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

7. Replace the left end cap.

Step 2: Power on the controllers


1. Connect your laptop to the serial console port so you can monitor the boot sequence when power is applied to the controllers.

- a. Set the serial console port on the laptop to 115,200 baud with N-8-1.




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

- b. Connect the console cable to the laptop, and connect the serial console port on the controller using the console cable that came with your platform.
- c. Connect the laptop to the switch on the management subnet.



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



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

Complete storage system setup and configuration - AFF A70, AFF A90

After you have installed your hardware, use the ONTAP software to configure and manage your storage system.

Step 1: Gather cluster information

If you have not already done so, gather the information you will need to configure your cluster, such as your cluster management interface port and IP address.

Use the [cluster setup worksheet](#) to record the values that you need during the cluster setup process. If a default value is provided, you can use that value or else enter your own.

Step 2: Discover your cluster network


The discovery process enables you to discover your storage system controllers on the network.

Option 1: Network discovery is enabled

If you have network discovery enabled on your laptop, you can complete platform setup and configuration using automatic cluster discovery.

Steps

1. Connect your laptop to the management switch and access the network computers and devices.
2. Select an ONTAP icon listed to discover:



- a. Open File Explorer.
- b. Click **Network** in the left pane.
- c. Right-click and select **refresh**.
- d. Double-click either ONTAP icon and accept any certificates displayed on your screen.



XXXXX is the platform serial number for the target node.

System Manager opens.

Option 2: Network discovery is not enabled


If network discovery is not enabled on your laptop, complete the configuration and setup using the ONTAP command line interface (CLI) Cluster Setup wizard.

Before you begin

Make sure your laptop is connected to the serial console port and the controllers are powered on. See [Power on the controllers](#) for instructions.

Steps


Assign an initial node management IP address to one of the nodes.

If the management network has DHCP...	Then...
Configured	Record the IP address assigned to the new controllers.
Not configured	<p>a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment.</p> <div style="display: flex; align-items: center; margin: 10px 0;">  <p>Check your laptop or console's online help if you do not know how to configure PuTTY.</p> </div> <p>b. Connect to the console of the first node.</p> <p>The node boots, and then the Cluster Setup wizard starts on the console.</p> <p>c. Enter the node's management IP address when prompted by the Cluster Setup wizard.</p>

Step 3: Configure your cluster

NetApp recommends that you use System Manager to set up new clusters. See [Configure ONTAP on a new cluster with System Manager](#) for setup instructions.

System Manager provides a simple and easy workflow for cluster set up and configuration including assigning a node management IP address, initializing the cluster, creating a local tier, configuring protocols and initial provisioning of attached storage.

 After your cluster is initialized, download and run [Active IQ Config Advisor](#) to confirm your setup.

Maintain

Maintain AFF A70, AFF A90 hardware

For the AFF A70, AFF A90 storage systems, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

Fan

The 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 NVRAM 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

Overview of boot media replacement - AFF A70, AFF A90

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots. Depending on your network configuration, you can perform either a nondisruptive or disruptive replacement.

You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_XXX.tgz`.

You also must copy the `image_XXX.tgz` file to the USB flash drive for later use in this procedure.

- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct controller:
 - The *impaired* controller is the controller on which you are performing maintenance.
 - The *healthy* controller is the HA partner of the impaired controller.

Check onboard encryption keys - AFF A70, AFF A90

Prior to shutting down the impaired controller and checking the status of the onboard encryption keys, you must check the status of the impaired controller, disable automatic giveback, and check the version of ONTAP that is running.

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](#).

Check NVE or NSE on systems running ONTAP 9.15 and later

Before shutting down the impaired controller, you need to verify whether the system has security key manager enabled or encrypted disks.

Verify security key-manager configuration


Steps

1. Determine if Key Manager is active with the `security key-manager keystore show` command. For more information, see the [security key-manager keystore show MAN page](#)



You may have additional key manager types. The types are `KMIP`, `AKV`, and `GCP`. The process for confirming these types is the same as confirming `external` or `onboard` key manager types.

- If no output is displayed, go to [shutdown the impaired controller](#) to shutdown the impaired node.
 - If the command displays output, the system has `security key-manager` active and you need to display the `Key Manager` type and status.
2. Display the information for the active `Key Manager` using the `security key-manager key query` command.
 - If the `Key Manager` type displays `external` and the `Restored` column displays `true`, it's safe to shut down the impaired controller.
 - If the `Key Manager` type displays `onboard` and the `Restored` column displays `true`, you need to complete some additional steps.
 - If the `Key Manager` type displays `external` and the `Restored` column displays anything other than `true`, you need to complete some additional steps.
 - If the `Key Manager` type displays `onboard` and the `Restored` column displays anything other than `true`, you need to complete some additional steps.
 3. If the `Key Manager` type displays `onboard` and the `Restored` column displays `true`, manually back up the OKM information:

- a. Enter `y` when prompted to continue: `set -priv advanced`
 - b. Enter the command to display the key management information: `security key-manager onboard show-backup`
 - c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM.
 - d. You can safely shut down the impaired controller.
4. If the `Key Manager` type displays `onboard` and the `Restored` column displays anything other than `true`:
- a. Enter the onboard security key-manager sync command: `security key-manager onboard sync`
- 

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

Shut down the controller - AFF A70, AFF A90

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller. Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:>`

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

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`



When you see *Do you want to disable auto-giveback?*, enter *y*.

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></code> When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <i>y</i> .

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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

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


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

Replace the boot media - AFF A70, AFF A90

To replace the boot media, you must remove the System Management module from the back of the system, remove the impaired boot media, install the replacement boot media in the System Management module, and transfer the boot image to a USB flash drive.


Step 1: Replace the boot media

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



<p>1</p>	<p>System Management module cam latch</p>
<p>2</p>	<p>Boot media locking button</p>
<p>3</p>	<p>Boot media</p>

1. On the front of the chassis, use your thumbs to firmly push each drive in until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.



- 2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
- 3. Disconnect power by pulling the controller module out about three inches:



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

- a. Press down on both of the locking latches, and then rotate both latches downward at the same time.
 - b. Pull the controller module about 3 inches out of the chassis to disengage power.
 - c. 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.
 - 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 system management cam button. The cam lever moves away from the chassis.
 - f. Rotate the cam lever all the way down and remove the System Management module from the controller module.
 - g. Place the System Management module on an anti-static mat, so that the boot media is accessible.
4. Remove the boot media from the management module:
 - a. Press the blue locking button.
 - b. Rotate the boot media up, slide it out of the socket, and set it aside.
 5. Install the replacement boot media into the System Management module:
 - a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
 - b. Rotate the boot media down toward the locking button.
 - c. Push the locking button, rotate the boot media all the way down and then release the locking button.
 6. Reinstall the System Management module:
 - a. Rotate the cable management tray up to the closed position.
 - b. Recable the System Management module.

Step 2: Transfer the boot image to the boot media

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

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- You must have a copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the [Downloads](#) section on the NetApp Support Site
 - If NVE is supported, download the image with NetApp Volume Encryption, as indicated in the download button.
 - If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection between the node management ports of the controllers (typically the e0M interfaces).

Steps

1. Download and copy the appropriate service image from the [NetApp Support Site](#) to the USB flash drive.

- a. Download the service image from the Downloads link on the page, to your work space on your laptop.
- b. Unzip the service image.



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

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

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

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

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

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

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

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

5. Set your network connection type at the LOADER prompt:

- If you are configuring DHCP: `ifconfig e0M -auto`



The target port you configure is the target port you use to communicate with the impaired controller from the healthy controller during var file system restore with a network connection. You can also use the e0M port in this command.

- If you are configuring manual connections: `ifconfig e0M -addr=filer_addr -mask=netmask -gw=gateway`
 - `filer_addr` is the IP address of the storage system.
 - `netmask` is the network mask of the management network that is connected to the HA partner.
 - `gateway` is the gateway for the network.



Other parameters might be necessary for your interface. You can enter `help ifconfig` at the firmware prompt for details.

Boot the recovery image - AFF A70, AFF A90


You must boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

Steps

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

The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

If your system has...	Then...
A network connection	<p>a. On the impaired controller, press Y when you see <code>Do you want to restore the backup configuration now?</code></p> <p>b. On the impaired controller, press Y when prompted to overwrite <code>/etc/ssh/ssh_host_ecdsa_key</code>.</p> <p>c. On the healthy partner controller, set the impaired controller to advanced privilege level: <code>set -privilege advanced</code>.</p> <p>d. On the healthy partner controller, run the restore backup command: <code>system node restore-backup -node local -target -address impaired_node_IP_address</code>.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin: 10px 0;">  If you see any message other than a successful restore, contact NetApp Support. </div> <p>e. On the healthy partner controller, return the impaired controller to admin level: <code>set -privilege admin</code>.</p> <p>f. On the impaired controller, press y when you see <code>Was the restore backup procedure successful?</code>.</p> <p>g. On the impaired controller, press y when you see <code>...would you like to use this restored copy now?</code>.</p> <p>h. On the impaired controller, press y when prompted to reboot the impaired controller and press <code>ctrl-c</code> for the Boot Menu.</p> <p>i. If the system does not use encryption, select <i>Option 1 Normal Boot.</i>, otherwise go to Restore key managers.</p> <p>j. Connect the console cable to the partner controller.</p> <p>k. Give back the controller using the <code>storage failover giveback -fromnode local</code> command.</p> <ol style="list-style-type: none"> 1. Restore automatic giveback if you disabled it by using the <code>storage failover modify -node local -auto-giveback true</code> command. <p>l. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the <code>system node autosupport invoke -node * -type all -message MAINT=END</code> command.</p>
No network connection	Contact NetApp Support .
No network connection and is in a MetroCluster IP configuration	Contact NetApp Support .

Encryption restore - AFF A70, AFF A90

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



If NSE or NVE are enabled along with Onboard or external Key Manager you must restore settings you captured at the beginning of this procedure.

Steps

1. Connect the console cable to the target controller.

Option 1: Systems with onboard key manager server configuration

Restore the onboard key manager configuration from the ONATP boot menu.

Before you begin

You need the following information while restoring the OKM configuration:

- Cluster-wide passphrase entered [while enabling onboard key management](#).
- [Backup information for the Onboard Key Manager](#).
- Perform the [How to verify onboard key management backup and cluster-wide passphrase](#) procedure before proceeding.

Steps

1. From the ONTAP boot menu select option 10:

```
Please choose one of the following:
```

```
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? _10_
```

2. Confirm the continuation of the process. This option must be used only in disaster recovery procedures. Are you sure? (y or n): **y**
3. Enter the cluster-wide passphrase twice.



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

```
Enter the passphrase for onboard key management:
```

```
Enter the passphrase again to confirm:
```

4. Enter the backup information. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

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


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

```
Successfully recovered keymanager secrets.
```

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



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

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

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

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

7. Confirm that the controller's console displays Waiting for giveback...(Press Ctrl-C to abort wait)

8. From the partner node, giveback the partner controller: `storage failover giveback -fromnode local -only-cfo-aggregates true`
9. Once booted only with CFO aggregate run the `security key-manager onboard sync` command:
10. Enter the cluster-wide passphrase for the Onboard Key Manager:

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

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

11. Ensure that all keys are synced: `security key-manager key query -restored false`

There are no entries matching your query.



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

12. Giveback of the node from the partner: `storage failover giveback -fromnode local`

Option 2: Systems with external key manager server configuration

Restore the external key manager configuration from the ONATP boot menu.

Before you begin

You need the following information for restoring the external key manager (EKM) configuration:

- You need a copy of the `/cfcard/kmip/servers.cfg` file from another cluster node, or, the following information:
- The KMIP server address.
- The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node, or, the client certificate.
- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node, or, the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node, or, the KMIP server CA(s).

Steps

1. Select Option 11 from the ONTAP boot menu.

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

2. When prompted confirm you have gathered the required information:

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

You may also these prompts instead:

- e. Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n} *n*
 - i. Do you know the KMIP server address? {y/n} *y*
 - ii. Do you know the KMIP Port? {y/n} *y*

3. Supply the information for each of these prompts:

- a. Enter the client certificate (client.crt) file contents:
- b. Enter the client key (client.key) file contents:
- c. Enter the KMIP server CA(s) (CA.pem) file contents:
- d. Enter the server configuration (servers.cfg) file contents:

Example

Enter the client certificate (client.crt) file contents:

```
-----BEGIN CERTIFICATE-----
MIIDvjCCAqagAwIBAgICN3gwDQYJKoZIhvcNAQELBQAwwY8xCzAJBgNVBAYTA1VT
MRMwEQYDVQQIEwpDYWxpZm9ybmlhMQwwCgYDVQQHEwNTVkwxDzANBgNVBAoTBk51
MSUubQusvzAFs8G3P54GG32iIRvaCFnj2gQpCxcilJ0qB2foiBGx5XVQ/Mtk+rlap
Pk4ECW/wqSOUXDYtJs1+RB+w0+SHx8mzxp bz3mXF/X/1PC3YOzVNCq5eieek62si
Fp8=
-----END CERTIFICATE-----
```

Enter the client key (client.key) file contents:

```
-----BEGIN RSA PRIVATE KEY-----
MIIEpQIBAAKCAQEAOUleaajEG6QC2h2Zih0jEaGVtQUexNeoCFwKPomSePmjDNtrU
MSB1SlX3VgCuElHk57XPdq6xSbYl b kIb4bAgLztHEmUDOkGmXYAkblQ=
-----END RSA PRIVATE KEY-----
```

Enter the KMIP server CA(s) (CA.pem) file contents:

```
-----BEGIN CERTIFICATE-----
MIIEIzCCA3OgAwIBAgIBADANBgkqhkiG9w0BAQsFADCBjzELMAkGA1UEBhMVCVMx
7yaumMQETNrpMfP+nQMd34y4AmseWYGM6qG0z37BRnYU0Wf2qDL61cQ3/jkm7Y94
EQBKG1NY8dVyjphmYZv+
-----END CERTIFICATE-----
```

Enter the IP address for the KMIP server: 10.10.10.10

Enter the port for the KMIP server [5696]:

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

4. The recovery process will complete:

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
[Aug 29 21:06:28]: 0x808806100: 0: DEBUG: kmp2::main:
[initOpenssl]:460: Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

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

```

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

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

```

Complete the boot media replacement

Complete the boot media replacement process after the normal boot by gcompleting final checks and giving back storage.

1. Check the console output:

If the console displays...	Then...
The login prompt	Go to Step 6.
Waiting for giveback...	<ul style="list-style-type: none"> a. Log into the partner controller. b. Confirm the target controller is ready for giveback with the <i>storage failover show</i> command.

2. Move the console cable to the partner controller and give back the target controller storage using the *storage failover giveback -fromnode local -only-cfo-aggregates true* command.
- If the command fails because of a failed disk, physically disengage the failed disk, but leave the disk in the slot until a replacement is received.
 - If the command fails because the partner is "not ready", wait 5 minutes for the HA subsystem to synchronize between the partners.

- If the command fails because of an NDMP, SnapMirror, or SnapVault process, disable the process. See the appropriate Documentation Center for more information.
3. Wait 3 minutes and check the failover status with the `storage failover show` command.
 4. At the clustershell prompt, enter the `network interface show -is-home false` command to list the logical interfaces that are not on their home controller and port.

If any interfaces are listed as `false`, revert those interfaces back to their home port using the `net int revert -vserver Cluster -lif _nodename` command.

5. Move the console cable to the target controller and run the `version -v` command to check the ONTAP versions.
6. Use the `storage encryption disk show` to review the output.
7. Use the `security key-manager key query` command to display the key IDs of the authentication keys that are stored on the key management servers.
 - If the `Restored` column = `yes/true`, you are done and can proceed to complete the replacement process.
 - If the `Key Manager type` = `external` and the `Restored` column = anything other than `yes/true`, use the `security key-manager external restore` command to restore the key IDs of the authentication keys.



If the command fails, contact Customer Support.

- If the `Key Manager type` = `onboard` and the `Restored` column = anything other than `yes/true`, use the `security key-manager onboard sync` command to synchronize the missing onboard keys on the repaired node.

Use the `security key-manager key query` command to verify that the `Restored` column = `yes/true` for all authentication keys.

8. Connect the console cable to the partner controller.
9. Give back the controller using the `storage failover giveback -fromnode local` command.
10. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
11. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

Return the failed part to NetApp - AFF A70, AFF A90

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

Chassis

Replace the chassis - AFF A70, AFF A90

To replace the chassis, you must move the bezel, controller modules, and NVMe drives from the impaired chassis to the replacement chassis, and then remove the impaired chassis from the equipment rack or system cabinet and install the replacement chassis in

its place.

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

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving the bezel, NVMe drives, and controller modules to the new chassis, and that the replacement chassis is a new component from NetApp.
- **This procedure is disruptive.** For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - AFF A70, AFF A90

This procedure is for 2-node, non-MetroCluster configurations only. If you have a system with more than two nodes, see [How to perform a graceful shutdown and power up of one HA pair in a 4-node cluster](#).

Before you begin

You need:

- Local administrator credentials for ONTAP.
- NetApp onboard key management (OKM) cluster-wide passphrase if using storage encryption or NVE/NAE.
- BMC accessibility for each controller.
- Stop all clients/host from accessing data on the NetApp system.
- Suspend external backup jobs.
- Necessary tools and equipment for the replacement.



If the system is a NetApp StorageGRID or ONTAP S3 used as FabricPool cloud tier, refer to the [Gracefully shutdown and power up your storage system Resolution Guide](#) after performing this procedure.



If using SSDs, refer to [SU490: \(Impact: Critical\) SSD Best Practices: Avoid risk of drive failure and data loss if powered off for more than two months](#)

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. Turn off AutoSupport and indicate how long you expect the system to be offline:

```
system node autosupport invoke -node * -type all -message "MAINT=8h Power Maintenance"
```

3. Identify the SP/BMC address of all nodes:

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

4. Exit the cluster shell: `exit`
5. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step.

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



Open an SSH session to every SP/BMC connection so that you can monitor progress.

6. Halt the 2 nodes located in the impaired chassis:

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

7. Enter **y** for each controller in the cluster when you see *Warning: Are you sure you want to halt node "cluster <node-name> number"? {y|n}*:
8. Wait for each controller to halt and display the LOADER prompt.


Move and replace hardware - AFF A70, AFF A90

Move the power supplies, hard drives, and controller module from the impaired chassis to the new chassis, and swap out the impaired chassis with the new chassis of the same model as the impaired chassis.


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.

1. On the front of the chassis, use your thumbs to firmly push each drive in until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.



2. Check the amber NVRAM 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 NVRAM status LED is flashing, it could mean the controller module was not taken over or halted properly (uncommitted data). If the impaired controller module was not successfully taken over by the partner controller module, contact [NetApp Support](#) before continuing with this procedure.

The general behavior of the NVRAM status LED on the impaired controller module is as follows:

- The NVRAM status LED flashes when power is removed from the controller module and the storage system is in the "waiting for giveback" state, or the controller module is not taken over or halted properly (uncommitted data).

- The NVRAM status LED flashes when the controller module is removed from the chassis and could mean the controller module is not taken over or halted properly (uncommitted data). Confirm that the controller module has been cleanly takeover by the partner controller module or the impaired controller module shows `waiting for giveback`. Then, the flashing LED can be ignored (and the controller module can be removed from the chassis).

3. If you are not already grounded, properly ground yourself.
4. 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.

5. 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.

6. Remove the cable management device from the controller module.
7. 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	a Locking latch
2	Locking pin

8. 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.

9. Repeat these steps for the other controller module in the chassis.

Step 2: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new

chassis.

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

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



When removing a drive, always use two hands to support its weight.



Drives are fragile. Handle them as little as possible to prevent damage to them.

3. Keep track of what drive bay each drive was from and set the drives aside on a static-free cart or table.

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

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
4. Slide the chassis all the way into the equipment rack or system cabinet.
5. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
6. Install the drives from the old chassis into the replacement chassis:
 - a. Align the drive from the old chassis with the same bay opening in the new chassis.
7. Gently push the drive into the chassis as far as it will go.

The cam handle engages and begins to rotate upward.

- a. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up and against the drive carrier.

Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It clicks when it is secure.

- b. Repeat the process for the remaining drives in the system.
8. If you have not already done so, install the bezel.

Step 4: Reinstall the controller modules

Reinstall the controller module and reboot 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 chassis, and then gently push the controller module halfway into the system.



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

3. Recable the system, as needed.

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

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.

- a. Rotate the locking latches upward into the locked position.

5. 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.

The controller module boots when power is restored. If it boots to the LOADER prompt, reboot the controller with the `boot_ontap` command.

6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
8. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - AFF A70, AFF A90

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

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

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:

a. Set the HA state for the chassis: `ha-config modify chassis HA-state`

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

- `ha`
- `mcc`

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

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

Step 2: 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 & Replacements](#) page for further information.

Controller

Overview of controller module replacement - AFF A70, AFF A90

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section [Choosing the correct recovery procedure](#) to determine whether you should use this procedure.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- 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 - AFF A70, AFF A90

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:>`

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

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`



When you see *Do you want to disable auto-giveback?*, enter *y*.

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></code> When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <i>y</i> .

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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

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

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


Replace the controller module hardware - AFF A70, AFF A90

To replace the controller, you must remove the impaired controller, move FRU components from the impaired controller module to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.


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.

1. On the front of the chassis, use your thumbs to firmly push each drive in until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.



2. 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 NVRAM status LED is flashing, it could mean the controller module was not taken over or halted properly (uncommitted data). If the impaired controller module was not successfully taken over by the partner controller module, contact [NetApp Support](#) before continuing with this procedure.

The general behavior of the NVRAM status LED on the impaired controller module is as follows:

- The NVRAM status LED flashes when power is removed from the controller module and the storage system is in the "waiting for giveback" state, or the controller module is not taken over or halted properly (uncommitted data).
- The NVRAM status LED flashes when the controller module is removed from the chassis and could mean the controller module is not taken over or halted properly (uncommitted data). Confirm that the controller module has been cleanly takeover by the partner controller module or the impaired controller module shows `waiting for giveback`. Then, the flashing LED can be ignored (and the controller module can be removed from the chassis).

3. If you are not already grounded, properly ground yourself.
4. 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.

5. 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.

6. Remove the cable management device from the controller module.
7. 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	a Locking latch
2	Locking pin

8. 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.

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.

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 the fan module by aligning its edges with the opening in the controller module, and then sliding the fan module into the controller module 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.

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.

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.



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.



1

I/O module cam lever

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.

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

a. Depress the target I/O module cam latch button.

The cam lever moves away from the controller module.

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.




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

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



1	Cam locking button
---	--------------------

2

DIMM locking tab

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

Step 9: Install the controller module

Reinstall the controller module and reboot 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 chassis, and then gently push the controller module halfway into the system.



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

3. Recable the system, as needed.

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

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.

- a. Rotate the locking latches upward into the locked position.
5. 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.

The controller module boots when power is restored. If it boots to the LOADER prompt, reboot the

controller with the `boot_ontap` command.

6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
8. If you have not already done so, reinstall the cable management device and recable the controller.

Restore and verify the system configuration - AFF A70, AFF A90

After completing the hardware replacement, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

The date and time are given in GMT.

Step 2: Verify and set the HA state of the chassis

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

1. In Maintenance mode from the new controller module, verify that all components display the same HA state: `ha-config show`

The HA state should be the same for all components.

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

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

- ha
- mcc (not supported)
- mccip
- non-ha (not supported)

3. If the displayed system state of the controller module does not match your system configuration, set the HA state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable and give back the controller - AFF A70, AFF A90

Continue the replacement procedure by recabling the storage, network connections and giving back the controller.

Step 1: Recable the system

Recable the controller module's storage and network connections.

Steps

1. Recable the the controller module to storage and network connections.
2. Verify that the cabling is correct by using [Active IQ Config Advisor](#).
 - a. Download and install Config Advisor.
 - b. Enter the information for the target system, and then click Collect Data.
 - c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
 - d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Give back the controller

1. If your storage system has Encryption configured, you must restore Storage or Volume Encryption functionality using the following procedure to reboot the system:
 - a. Boot to Menu and run Option 10
 - b. Input the passphrase & backup up data, then do Normal boot see [Restore onboard key management encryption keys](#).
 - c. Perform CFO only giveback
 - d. Perform Onboard Sync and verify SVM-KEK is set to true see [Giveback after MB replacement fails - operation was vetoed by keymanager](#)
 - e. Giveback SFO, (no force)
2. If your system does not have Encryption configured, complete the following procedure to reboot the system:

- a. Boot to Menu and run Option 1.
- b. Give back the controller:
- c. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

The *replacement* controller takes back its storage and completes booting.



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

[Find the High-Availability Configuration content for your version of ONTAP 9](#)

- d. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`
3. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

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

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

This is required if both of the following are true:

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

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

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

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

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

4 entries were displayed.
```

6. Verify that the expected volumes are present for each controller: `vol show -node node-name`

7. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - AFF A70, AFF A90

To restore your system to full operation, you must verify the LIFs, check cluster health, and return the failed part to NetApp.

Step 1: Verify LIFs and check cluster health

Before returning the *replacement* node to service, you should verify that the LIFs 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 LIFs 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 for more information.
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 & Replacements](#) page for further information.

Replace a DIMM - AFF A70, AFF A90

You must replace a DIMM in the controller module when your system registers an increasing number of correctable error correction codes (ECC); failure to do so causes a system panic.

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 using one of the following options.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:>`

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

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`



When you see *Do you want to disable auto-giveback?*, enter *y*.

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></code> When the impaired controller shows <i>Waiting for giveback...</i> , press Ctrl-C, and then respond <i>y</i> .

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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


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

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


Step 2: 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.

1. On the front of the chassis, use your thumbs to firmly push each drive in until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.



2. Check the amber NVRAM 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 NVRAM status LED is flashing, it could mean the controller module was not taken over or halted properly (uncommitted data). If the impaired controller module was not successfully taken over by the partner controller module, contact [NetApp Support](#) before continuing with this procedure.

The general behavior of the NVRAM status LED on the impaired controller module is as follows:

- The NVRAM status LED flashes when power is removed from the controller module and the storage system is in the "waiting for giveback" state, or the controller module is not taken over or halted properly (uncommitted data).

- The NVRAM status LED flashes when the controller module is removed from the chassis and could mean the controller module is not taken over or halted properly (uncommitted data). Confirm that the controller module has been cleanly takeover by the partner controller module or the impaired controller module shows `waiting for giveback`. Then, the flashing LED can be ignored (and the controller module can be removed from the chassis).

3. If you are not already grounded, properly ground yourself.
4. 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.

5. 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.

6. Remove the cable management device from the controller module.
7. 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	a Locking latch
2	Locking pin

8. 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.

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.




Consult either the [Netapp Hardware Universe](#) or the FRU map on your controller module for exact DIMM locations for the AFF A70 or AFF A90.

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

1. 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.

2. 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.

3. 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.
4. Close the controller air duct.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot 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 chassis, and then gently push the controller module halfway into the system.



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

3. Recable the system, as needed.

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

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.

- a. Rotate the locking latches upward into the locked position.

5. 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.

The controller module boots when power is restored. If it boots to the LOADER prompt, reboot the controller with the `boot_ontap` command.

6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

Step 5: Return the failed part to NetApp

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

Replace SSD Drive - AFF A70, AFF A90

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 the [NetApp Hardware Universe](#).
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

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

When replacing 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.

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled



You manually assign drive ownership and then reenables automatic drive assignment later in this procedure.

- a. Verify whether automatic drive assignment is enabled: `storage disk option show`

You can enter the command on either controller module.

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

- b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

You must disable automatic drive assignment on both controller modules.

2. Properly ground yourself.
3. Physically identify the failed drive.

When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.



The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.

4. Remove the failed drive:
 - a. Press the release button on the drive face to open the cam handle.
 - b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

This allows the system to recognize that a drive was removed.

6. Insert the replacement drive:
 - a. With the cam handle in the open position, use both hands to insert the replacement drive.
 - b. Push until the drive stops.
 - c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

7. Verify that the drive's activity (green) LED is illuminated.

When the drive's activity LED is solid, it means that the drive has power. When the drive's activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat Steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenables automatic drive assignment if needed.

- a. Display all unowned drives: `storage disk show -container-type unassigned`

You can enter the command on either controller module.

- b. Assign each drive: `storage disk assign -disk disk_name -owner node_name`

You can enter the command on either controller module.

You can use the wildcard character to assign more than one drive at once.

- c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

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

Contact technical support at [NetApp Support](#), 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace a fan module - AFF A70, AFF A90

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:>`

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

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`



When you see *Do you want to disable auto-giveback?*, enter *y*.

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></code> When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <i>y</i> .

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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


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

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


Step 2: 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.

1. On the front of the chassis, use your thumbs to firmly push each drive in until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.



2. 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 NVRAM status LED is flashing, it could mean the controller module was not taken over or halted properly (uncommitted data). If the impaired controller module was not successfully taken over by the partner controller module, contact [NetApp Support](#) before continuing with this procedure.

The general behavior of the NVRAM status LED on the impaired controller module is as follows:

- The NVRAM status LED flashes when power is removed from the controller module and the storage system is in the "waiting for giveback" state, or the controller module is not taken over or halted properly (uncommitted data).

- The NVRAM status LED flashes when the controller module is removed from the chassis and could mean the controller module is not taken over or halted properly (uncommitted data). Confirm that the controller module has been cleanly takeover by the partner controller module or the impaired controller module shows `waiting for giveback`. Then, the flashing LED can be ignored (and the controller module can be removed from the chassis).

3. If you are not already grounded, properly ground yourself.
4. 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.

5. 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.

6. Remove the cable management device from the controller module.
7. 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	a Locking latch
2	Locking pin


8. 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 fan

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

1. Identify the fan module that you must replace by checking the console error messages or by locating the lit LED for the fan module on the motherboard.
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.

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 system, as needed.

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

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.

a. Rotate the locking latches upward into the locked position.

5. 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.

The controller module boots when power is restored. If it boots to the LOADER prompt, reboot the controller with the `boot_ontap` command.

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

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

Step 5: Return the failed part to NetApp

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

Replace NVRAM - AFF A70, AFF A90

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. To replace a failed NVRAM module, you must remove the module from the chassis, move the DIMMs to the replacement module, and install the replacement NVRAM module into the chassis.

All other components in the system must be functioning properly; if not, you must contact [NetApp 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 using one of the following options.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:>`

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

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`



When you see *Do you want to disable auto-giveback?*, enter *y*.

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></code> When the impaired controller shows <i>Waiting for giveback...</i> , press Ctrl-C, and then respond <i>y</i> .

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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


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

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

Step 2: 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.

1. On the front of the chassis, use your thumbs to firmly push each drive in until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.




2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Unplug the power cord from both PSUs.



If your system has DC power supplies, disconnect the power cable block from the PSU.

4. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
5. 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

6. Set the NVRAM module on a stable surface.
7. Remove the DIMMs, one at a time, from the impaired NVRAM module and install them in the replacement NVRAM module.
8. 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.
9. Recable the PSU.


 If your system has DC power supplies, reconnect the power cable block to the PSU.

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

Step 3: Replace a NVRAM DIMM

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


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

 If your system has DC power supplies, disconnect the power cable block from the PSU.

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 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

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 chassis:
 - a. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then

rotate the cam latch all the way up to lock the module in place.

10. Recable the PSU.



If your system has DC power supplies, reconnect the power cable block to the PSU.

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

Step 4: Reboot the controller

After you replace the FRU, you must reboot the controller module.

1. To boot ONTAP from the LOADER prompt, enter *bye*.

Step 5: Reassign disks

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



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

Steps

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

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

```
node1:> storage failover show
```

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

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

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

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

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

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

4 entries were displayed.
```

9. Verify that the expected volumes are present for each controller: *vol show -node node-name*
10. If storage encryption is enabled, you must restore functionality.
11. If you disabled automatic takeover on reboot, enable it from the healthy controller: *storage failover modify -node replacement-node-name -onreboot true*
12. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the *system node autosupport invoke -node * -type all -message MAINT=END* command.

Step 6: Return the failed part to NetApp

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

Replace the NV battery - AFF A70, AFF A90

To replace the NV battery, you must remove the controller module, remove the battery, replace the battery, and then reinstall the controller module.

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:>`

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

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`



When you see *Do you want to disable auto-giveback?*, enter *y*.

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></code> When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <i>y</i> .

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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


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

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


Step 2: 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.

1. On the front of the chassis, use your thumbs to firmly push each drive in until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.



2. 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 NVRAM status LED is flashing, it could mean the controller module was not taken over or halted properly (uncommitted data). If the impaired controller module was not successfully taken over by the partner controller module, contact [NetApp Support](#) before continuing with this procedure.

The general behavior of the NVRAM status LED on the impaired controller module is as follows:

- The NVRAM status LED flashes when power is removed from the controller module and the storage system is in the "waiting for giveback" state, or the controller module is not taken over or halted properly (uncommitted data).

- The NVRAM status LED flashes when the controller module is removed from the chassis and could mean the controller module is not taken over or halted properly (uncommitted data). Confirm that the controller module has been cleanly takeover by the partner controller module or the impaired controller module shows `waiting for giveback`. Then, the flashing LED can be ignored (and the controller module can be removed from the chassis).

3. If you are not already grounded, properly ground yourself.
4. 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.

5. 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.

6. Remove the cable management device from the controller module.
7. 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	a Locking latch
2	Locking pin


8. 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.

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.

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 system, as needed.

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

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.

- a. Rotate the locking latches upward into the locked position.

5. 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.

The controller module boots when power is restored. If it boots to the LOADER prompt, reboot the controller with the `boot_ontap` command.

6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

Step 5: Return the failed part to NetApp

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

I/O module

Overview of add and replace I/O module - AFF A70-90K

You can replace a failed I/O module in your 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 I/O module - AFF A70, AFF A90

You can add an I/O module to your storage system by either adding a new I/O module into a storage system with empty slots or by replacing an I/O module with a new one in a fully-populated storage system.

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.

Option 1: Add an I/O module to a storage system with empty slots

You can add an I/O module into an empty module slot in your storage system.

Step 1: Shut down the impaired controller module

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

Option 1: Most systems

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

About this task

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

Steps

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

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

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

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

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: <code>storage failover takeover -ofnode <i>impaired_node_name</i></code> When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <i>y</i> .

Option 2: Controller is in a MetroCluster



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

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

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

Configuration State is configured and that the nodes are in an enabled and normal state (metrocluster node show).

Steps

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

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

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

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

2. Disable automatic giveback from the console of the healthy controller:

```
storage failover modify  
-node local -auto-giveback false
```
3. Take the impaired controller to the LOADER prompt:


If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next Step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: <pre>storage failover takeover -ofnode impaired_node_name</pre> When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <i>y</i> .

Step 2: Add I/O modules

1. On the front of the chassis, use your thumbs to firmly push each drive in until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.




Make sure NVRAM destage has completed before proceeding.



2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Disconnect power by pulling the controller module out about three inches:
 - a. Press down on both of the locking latches, and then rotate both latches downward at the same time.
 - b. Pull the controller module about 3 inches out of the chassis to disengage power.
4. 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.
5. Remove the target slot blanking module from the chassis:



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 on the blanking module in the target slot.

- 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 blanking module.
6. Install the I/O module:
- a. Align the I/O module with the edges of the controller module slot opening.
 - b. Gently slide the module all the way into the into the slot, and then rotate the cam latch all the way up to lock the module in place.
7. Cable the I/O module.

If the I/O module is a NIC, cable the module to the data switches.

If the I/O module is a storage module, cable it to the NS224 shelf.



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

8. Reinstall the controller module into the chassis.



If you have DC PSUs, reconnect the cable blocks to the PSUs.

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

10. Reboot the controller from the LOADER prompt: *bye*



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

11. Give back the controller from the partner controller: *storage failover giveback -ofnode target_node_name*
12. Repeat these steps for controller B.
13. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
14. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
15. If you installed a storage I/O module, install and cable your NS224 shelves, as described in [Hot-adds a shelf](#).

Option 2: Add an I/O module in a storage system with no empty slots

You can change an I/O module in an I/O slot in a fully-populated system by removing an existing I/O module and replacing it with a different I/O module.

1. If you are:

Replacing a...	Then...
NIC I/O module with the same the same number of ports	The LIFs will automatically migrate when its controller module is shut down.

Replacing a...	Then...
NIC I/O module with fewer ports	Permanently reassign the ASAected LIFs to a different home port. See Migrating a LIF for information about using System Manager to permanently move the LIFs.
NIC I/O module with a storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF .

Step 1: Shut down the impaired controller module

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

Option 1: Most systems

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

About this task

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

Steps

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

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

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

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

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <code>y</code> when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: <code>storage failover takeover -ofnode <i>impaired_node_name</i></code> When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <code>y</code> .

Option 2: Controller is in a MetroCluster



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

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

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

Configuration State is configured and that the nodes are in an enabled and normal state (metrocluster node show).

Steps

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

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

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

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

2. Disable automatic giveback from the console of the healthy controller:


```
storage failover modify  
-node local -auto-giveback false
```
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next Step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: <pre>storage failover takeover -ofnode impaired_node_name</pre> When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <i>y</i> .

Step 2: Replace an I/O module

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

1. On the front of the chassis, use your thumbs to firmly push each drive in until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.



2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Disconnect power by pulling the controller module out about three inches:




If your storage system has DC power supplies, disconnect the power cable block from the PSU.

- a. Press down on both of the locking latches, and then rotate both latches downward at the same time.
 - b. Pull the controller module about 3 inches out of the chassis to disengage power.
4. Unplug any cabling on the target I/O module.
 5. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
 6. Remove the target I/O module from the chassis:



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.

The cam latch moves away from the chassis.

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 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.

7. Install the 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 chassis, and then rotate the cam latch all the way up to lock the module in place.

8. Cable the I/O module.

9. Repeat the remove and install steps to replace additional modules for the controller module.

10. Reinstall the controller module into the chassis.



If you have DC PSUs, reconnect the cable blocks to the PSUs.

11. Rotate the cable management tray into the locked position.

12. Reboot the controller module from the LOADER prompt: `_bye_`

a. Check the version of BMC on the controller: `system service-processor show`

b. Update the BMC firmware if needed: `system service-processor image update`

c. Reboot the node: `bye`



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



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

13. Give back the controller module from the partner controller module. `storage failover giveback -ofnode target_node_name`

14. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto-giveback true`

15. If you added:

If I/O module is a...	Then...
NIC module	Use the <code>storage port modify -node *<node name> -port *<port name> -mode network</code> command for each port.

If I/O module is a...	Then...
Storage module	Install and cable your NS224 shelves, as described in Hot-add a shelf .

16. Repeat these steps for controller B.

Replace I/O module - AFF A70, AFF A90

Use this procedure to replace a failed I/O module.

- You can use this procedure with all versions of ONTAP supported by your storage system.
- All other components in the storage system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:>`

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

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`



When you see *Do you want to disable auto-giveback?*, enter *y*.

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></code> When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <i>y</i> .

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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


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

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

Step 2: Replace a failed I/O module

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

1. On the front of the chassis, use your thumbs to firmly push each drive in until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.



2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Disconnect power by pulling the controller module out about three inches:



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


- a. Press down on both of the locking latches, and then rotate both latches downward at the same time.
 - b. Pull the controller module about 3 inches out of the chassis to disengage power.
4. Unplug any cabling on the target I/O module.

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

5. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
6. 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.

The cam latch moves away from the controller module.

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.

7. Set the I/O module aside.

8. 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.

9. Cable the I/O module.

10. Repeat the remove and install steps to replace additional modules for the controller.

11. Reinstall the controller module into the chassis.



If you have DC power supplies, reconnect the cable blocks to the PSUs.

12. 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 module.



If the new I/O module is not the same model as the failed module, you must first reboot the BMC.

Steps

1. Reboot the BMC if the replacement module is not the same model as the old module:
 - a. From the LOADER prompt, change to advanced privilege mode: *set privilege advanced*
 - b. Reboot the BMC: *sp reboot*
2. From the LOADER prompt, reboot the node: *bye*



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



Be sure to exit Maintenance mode after completing the conversion.

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

Step 4: Return the failed part to NetApp

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

Replace a power supply - AFF A70, AFF A90

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting it to the power source.

The power supplies are redundant and hot-swappable.

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: Replace 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 & Replacements](#) page for further information.

Option 2: Replace 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 & Replacements](#) page for further information.

Replace the real-time clock battery - AFF A70, AFF A90

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system.
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:>`

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

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`



When you see *Do you want to disable auto-giveback?*, enter *y*.

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></code> When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <i>y</i> .

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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


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

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


Step 2: 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.

1. On the front of the chassis, use your thumbs to firmly push each drive in until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.



2. 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 NVRAM status LED is flashing, it could mean the controller module was not taken over or halted properly (uncommitted data). If the impaired controller module was not successfully taken over by the partner controller module, contact [NetApp Support](#) before continuing with this procedure.

The general behavior of the NVRAM status LED on the impaired controller module is as follows:

- The NVRAM status LED flashes when power is removed from the controller module and the storage system is in the "waiting for giveback" state, or the controller module is not taken over or halted properly (uncommitted data).

- The NVRAM status LED flashes when the controller module is removed from the chassis and could mean the controller module is not taken over or halted properly (uncommitted data). Confirm that the controller module has been cleanly takeover by the partner controller module or the impaired controller module shows `waiting for giveback`. Then, the flashing LED can be ignored (and the controller module can be removed from the chassis).

3. If you are not already grounded, properly ground yourself.
4. 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.

5. 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.

6. Remove the cable management device from the controller module.
7. 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	a Locking latch
2	Locking pin


8. 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.

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.

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 system, as needed.

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

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.

- a. Rotate the locking latches upward into the locked position.

5. 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.

The controller module boots when power is restored. If it boots to the LOADER prompt, reboot the controller with the `boot_ontap` command.

6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

Step 5: Reset the time and date on the controller



After replacing the RTC battery, inserting controller and powering on first BIOS reset, you will see the following error messages: `RTC date/time error`. Reset date/time to default `RTC power failure error`. These messages are expected and you can continue with this procedure.

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



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

- a. At the LOADER prompt on the target controller, check the time and date with the `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.
 3. Return the 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`

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 & Replacements](#) page for further information.

Replace system management module - AFF A70, AFF A90

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

To replace the System Management module or the boot media, you must shut down the impaired controller.

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.

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:>`

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

2. Disable automatic giveback from the console of the healthy controller:

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



When you see *Do you want to disable auto-giveback?*, enter *y*.

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 impaired_node_name</pre> When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <i>y</i> .

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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

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

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


Step 2: Replace the impaired System Management module

Replace the impaired system management module.

1. On the front of the chassis, use your thumbs to firmly push each drive in until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.



Make sure NVRAM destage has completed before proceeding.




2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Disconnect power by pulling the controller module out about three inches:



If your system has DC power supplies, disconnect the power cable block from the PSU.


- a. Press down on both of the locking latches, and then rotate both latches downward at the same time.
 - b. Pull the controller module about 3 inches out of the chassis to disengage power.
4. 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.
 5. Remove the System Management module:
 - a. Remove any cables connected to the System Management module. Make sure that label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.



1

System Management module cam latch

6. 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.
7. 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.
8. 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.
 9. Recable the System Management module.



If your system has DC power supplies, reconnect the power cable block to the PSU.

10. Reinstall the controller module into the chassis.
11. Rotate the cable management tray up to the closed position.

Step 3: Reboot the controller module

Reboot the controller module.

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 by using the `storage failover modify -node local -auto -giveback true` command.
4. 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 for the node as soon as possible.

Before you begin

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

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

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 & Replacements](#) page for further information.

Copyright information

Copyright © 2024 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.