NetApp

ASA A900 systems

Install and maintain

NetApp
February 13, 2026

Table of Contents

ASA A900 systems

Install and setup

Start here: Choose your installation and setup experience

You can choose from different content formats to guide you through installing and setting up your new storage system.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

Quick steps - ASA A900

The quick guide provides graphic instructions for a typical installation of your system, from racking and cabling, through initial system bring-up. Use this content if you are familiar with installing NetApp systems.

Use the xref:./asa900/AFF A900 Installation and Setup Instructions

ⓘ | The ASA A900 uses the same installation procedure as the AFF A900 system.

Video steps - ASA A900

The following video shows how to install and cable your new system.

Animation - AFF A900 Installation and setup instructions

ⓘ | The ASA A900 uses the same installation procedure as the AFF A900 system.

Detailed steps - ASA 900

This page provides detailed step-by-step instructions for installing a typical NetApp system. Use this article if you want more detailed installation instructions.

Step 1: Prepare for installation

To install your system, you need to create an account on the NetApp Support Site, register your system, and get license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

You need to have access to the NetApp Hardware Universe for information about site requirements as well as additional information on your configured system.

What you need

You might also want to have access to the ONTAP 9 Release Notes for your version of ONTAP for more information about this system.

You need to provide the following at your site:

- Rack space for the storage system
- Phillips #2 screwdriver
- Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

Steps

1. Unpack the contents of all boxes.

2. Record the system serial number from the controllers.

> SSN: XXYYYYYYYYYY
> ‖‖‖‖‖‖‖‖‖‖‖‖‖‖‖‖

3. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the Hardware Universe to locate the cable and identify its use.

   NetApp Hardware Universe

Type of cable…	Part number and length	Connector type	For…
25 GbE data Cable	X66240A-05 (112-00639), 0.5m X66240A-2 (112-00598), 2m X66240A-5 (112-00600), 5m		Network cable
32 Gb FC (SFP+ Op)	X66250-2 (112-00342), 2m X66250-5 (112-00344), 5m X66250-15 (112-00346), 15m		FC optical network cable
40 GbE network cable	X66100-1 (112-00542), 1m X66100-3 (112-00543), 3m X66100-5 (112-00544), 5m		Ethernet data, cluster network

Type of cable…	Part number and length	Connector type	For…
100 GbE cable	X66211B-1 (112-00573), 1m X66211B-2 (112-00574), 2m X66211B-5 (112-00576), 5m		Network, NVME storage, Ethernet data, cluster network
Optical cables	X66031A (112-00436), 1m X66032A (112-00437), 2m X66033A (112-00438), 3m		FC optical network
Cat 6, RJ-45 (order dependent)	Part numbers X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network and Ethernet data
Micro-USB console cable	Not applicable		Console connection during software setup on non-Windows or Mac laptop/console
Power cables	Not applicable		Powering up the system

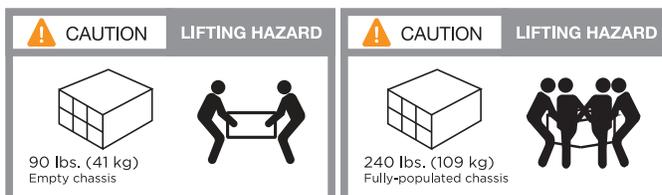4. Review the ONTAP Configuration Guide and collect the required information listed in that guide.

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

1. Install the rail kits, as needed.
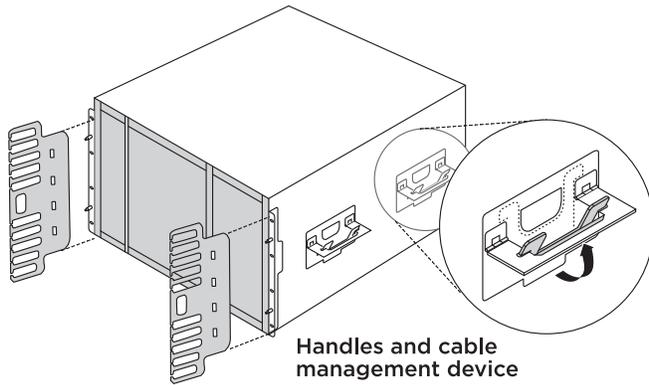
2. Install and secure your system using the instructions included with the rail kit.

   ⓘ You need to be aware of the safety concerns associated with the weight of the system.

⚠ CAUTION LIFTING HAZARD	⚠ CAUTION LIFTING HAZARD
90 lbs. (41 kg) Empty chassis	240 lbs. (109 kg) Fully-populated chassis

3. Attach cable management devices (as shown).

Handles and cable management device
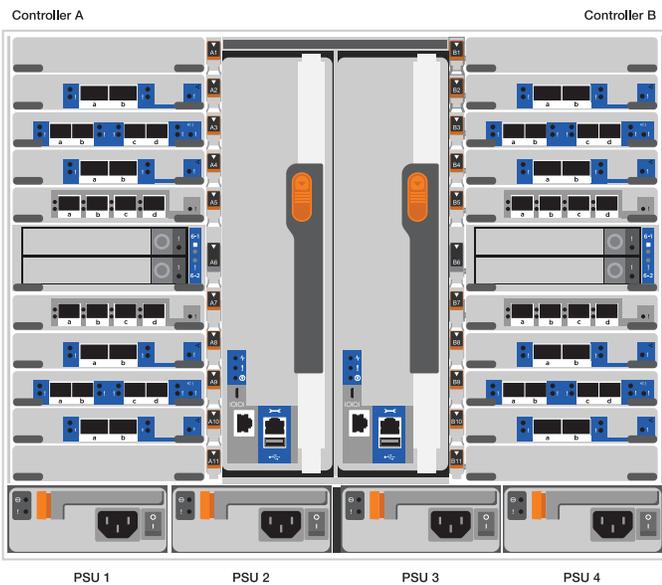
4. Place the bezel on the front of the system.

The following diagram shows a representation of what a typical system looks like and where the major components are located at the rear of the system:

Step 3: Cable controllers to your network

You can cable the controllers to your network by using the two-node switchless cluster method or by using the cluster interconnect network.

Option 1: Two-node switchless cluster

Management network, data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled on both controllers.

Before you begin

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all networking module ports.
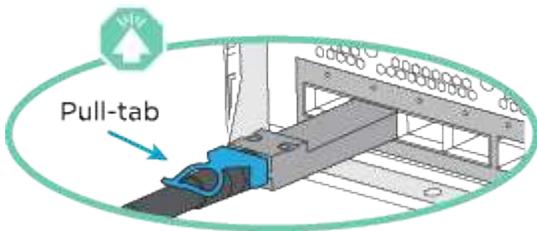
> ℹ️ As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

Animation - Cable a two-node switchless cluster

Step	Perform on each controller
1	Cable cluster interconnect ports: • Slot A4 and B4 (e4a) • Slot A8 and B8 (e8a)
2	Cable controller management (wrench) ports.
3	Cable 25 GbE network switches: Ports in slot A3 and B3 (e3a and e3c) and slot A9 and B9 (e9a and e9c) to the 25 GbE network switches. 40GbE host network switches: Cable host-side b ports in slot A4 and B4 (e4b) and slot A8 and B8 (e8b) to the host switch.
4	Cable 32 Gb FC connections: Cable ports in slot A5 and B5 (5a, 5b, 5c, and 5d) and slot A7 and B7 (7a, 7b, 7c, and 7d) to the 32 Gb FC network switches.
5	• Strap the cables to the cable management arms (not shown). • Connect the power cables to the PSUs and connect them to different power sources (not shown). PSU 1 and 3 provide power to all side A components, while PSU2 and PSU4 provide power to all side B components.

Option 2: Switched cluster

Management network, data network, and management ports on the controllers are connected to switches. The cluster interconnect and HA ports are cabled on to the cluster/HA switch.

Before you begin

You must have contacted your network administrator for information about connecting the system to the switches.

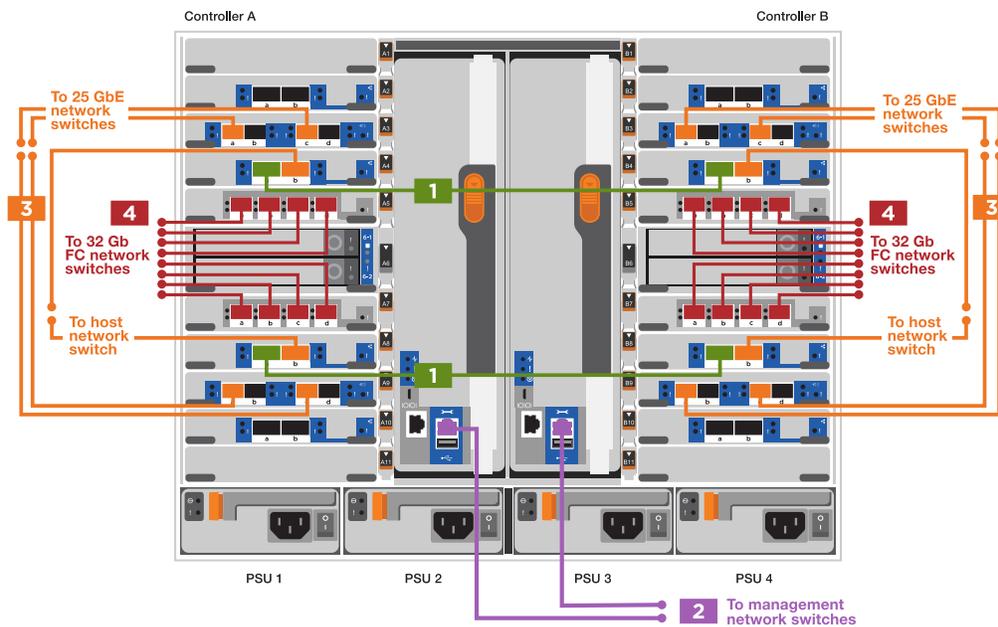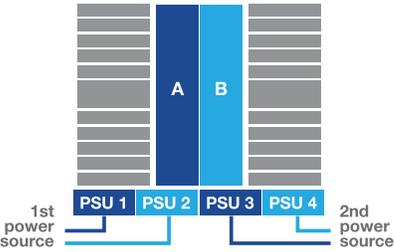Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all networking module ports.
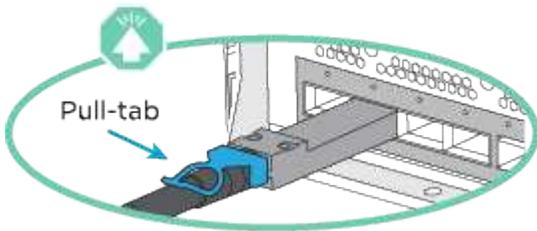
> ⓘ As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it over and try again.

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

Animation - Cable a switched cluster

Step	Perform on each controller
1	Cable cluster interconnect a ports: • Slot A4 and B4 (e4a) to the cluster network switch. • Slot A8 and B8 (e8a) to the cluster network switch.
2	Cable controller management (wrench) ports.
3	Cable 25GbE network switches: Ports in slot A3 and B3 (e3a and e3c) and slot A9 and B9 (e9a and e9c) to the 25 GbE network switches. 40GbE host network switches: Cable host-side b ports in slot A4 and B4 (e4b) and slot A8 and B8 (e8b) to the host switch.
4	Cable 32 Gb FC connections: Cable ports in slot A5 and B5 (5a, 5b, 5c, and 5d) and slot A7 and B7 (7a, 7b, 7c, and 7d) to the 32 Gb FC network switches.

| Step	Perform on each controller
5	• Strap the cables to the cable management arms (not shown). • Connect the power cables to the PSUs and connect them to different power sources (not shown). PSU 1 and 3 provide power to all side A components, while PSU2 and PSU4 provide power to all side B components. ![diagram of power cable and PSU connections showing PSU 1, PSU 2, PSU 3, PSU 4 with 1st power source and 2nd power source, and controllers A and B]
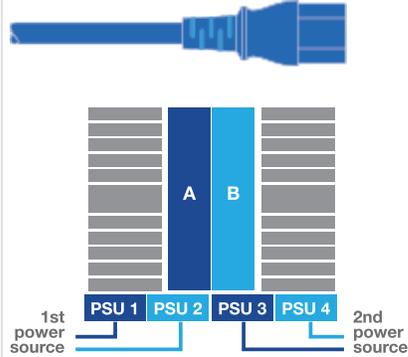
Step 4: Cable controllers to drive shelves

Cable either a single NS224 drive shelf or two NS224 drive shelves to your controllers.

Option 1: Cable the controllers to a single NS224 drive shelf

You must cable each controller to the NSM modules on the NS224 drive shelf.

Before you begin

- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the storage modules are up, while the pull tabs on the shelves are down.
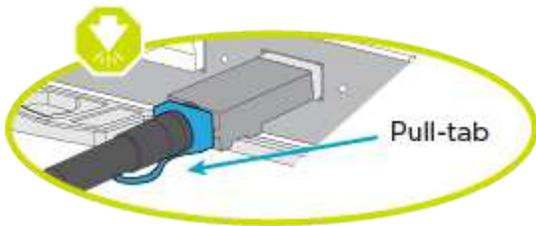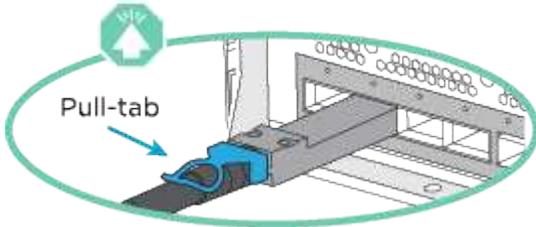
(i) As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

1. Use the following animation or drawings to cable your controllers to a single NS224 drive shelf.

Animation - Cable a single NS224 shelf

Step	Perform on each controller
1	• Connect controller A port e2a to port e0a on NSM A on the shelf. • Connect controller A port e10b to port e0b on NSM B on the shelf. 100 GbE cable
2	• Connect controller B port e2a to port e0a on NSM B on the shelf. • Connect controller B port e10b to port e0b on NSM A on the shelf. 100 GbE cable

Option 2: Cable the controllers to two NS224 drive shelves

You must cable each controller to the NSM modules on the NS224 drive shelves.

Before you begin

- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the storage modules are up, while the pull tabs on the shelves are down.
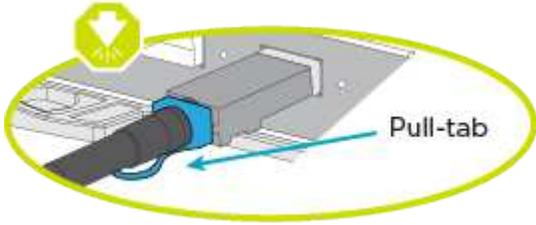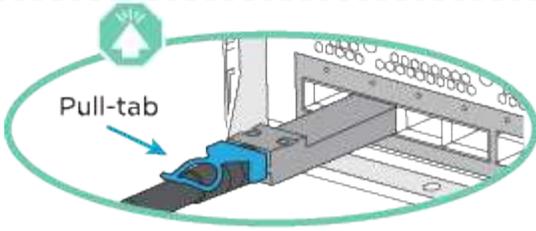
> ⓘ As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

1. Use the following animation or diagram to cable your controllers to two NS224 drive shelves.

   Animation - Cable two NS224 shelves

Step	Perform on each controller
1	• Connect controller A port e2a to NSM A e0a on shelf 1. • Connect controller A port e10b to NSM B e0b on shelf 1. • Connect controller A port e2b to NSM B e0b on shelf 2. • Connect controller A port e10a to NSM A e0a on shelf 2. 100 GbE cable
2	• Connect controller B port e2a to NSM B e0a on shelf 1. • Connect controller B port e10b to NSM A e0b on shelf 1. • Connect controller B port e2b to NSM A e0b on shelf 2. • Connect controller B port e10a to NSM B e0a on shelf 2. 100 GbE cable

Step 5: Complete system setup and configuration

You can complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: If network discovery is enabled

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

1. Use the following animation or drawing to set one or more drive shelf IDs:

   The NS224 shelves are pre-set to shelf ID 00 and 01. If you want to change the shelf IDs, you must create a tool to insert into the hole where button is located. see Change a shelf ID - NS224 shelves for detailed instructions.

   Animation - Set NVMe drive shelf IDs
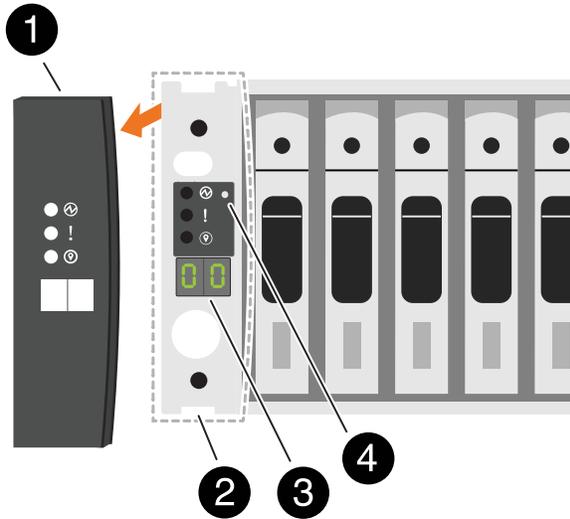
❶	Shelf end cap
❷	Shelf faceplate
❸	Shelf ID LED
❹	Shelf ID setting button

2. Turn on the power switches on the power supplies to both nodes.

   Animation - Turn on the power to the controllers

| | PSU 1 | PSU 2 | PSU 3 | PSU 4 |
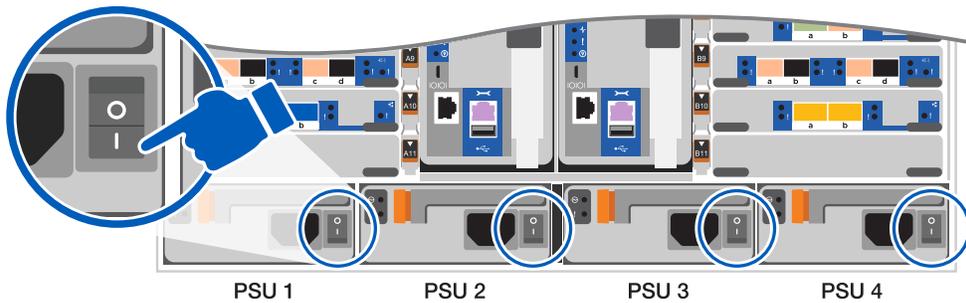
> ℹ️   Initial booting may take up to eight minutes.

3. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

4. Use the following animation to connect your laptop to the Management switch.

   [Animation - Connect your laptop to the Management switch](#)

5. Select an ONTAP icon listed to discover:
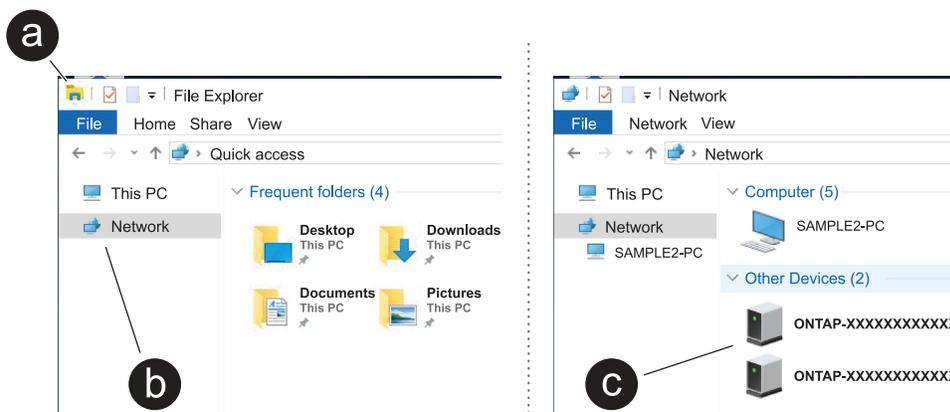
a. Open File Explorer.

b. Click **Network** in the left pane and right-click and select **refresh**.

c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

   > ℹ️   XXXXX is the system serial number for the target node.

   System Manager opens.

6. Use System Manager guided setup to configure your system using the data you collected in the ONTAP Configuration Guide.

7. Set up your account and download Active IQ Config Advisor:

   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

8. Verify the health of your system by running Config Advisor.

9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: If network discovery is not enabled

If you are not using a Windows or Mac-based laptop or console or if auto discovery is not enabled, you must complete the configuration and setup using this task.

1. Cable and configure your laptop or console:

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

      > ⓘ See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console using the console cable that came with your system, and then connect the laptop to the management switch on the management subnet.
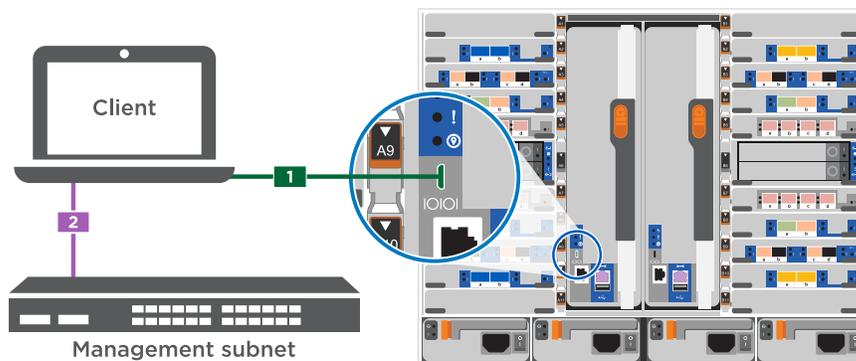
   c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.

2. Use the following animation to set one or more drive shelf IDs:

   The NS224 shelves are pre-set to shelf ID 00 and 01. If you want to change the shelf IDs, you must create a tool to insert into the hole where button is located. see Change a shelf ID - NS224 shelves for detailed instructions.

   Animation - Set NVMe drive shelf IDs
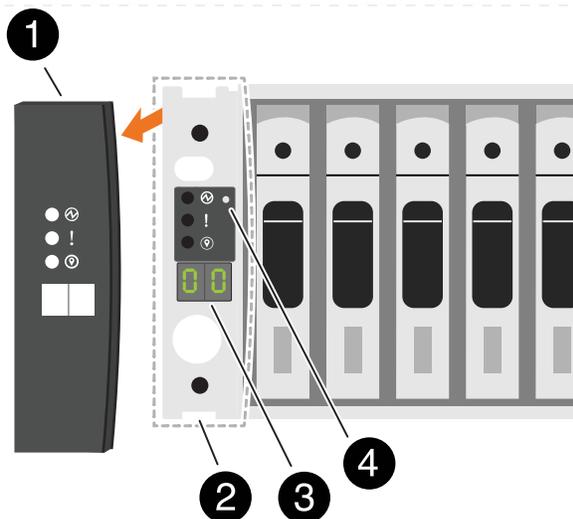
❶	Shelf end cap
❷	Shelf faceplate
❸	Shelf ID LED
❹	Shelf ID setting button

3. Turn on the power switches on the power supplies to both nodes.

   [Animation - Turn on the power to the controllers]

   image:[Callout number 1] drw_a900_power-on_IEOPS-941.svg[width=500px]

   > ⓘ Initial booting may take up to eight minutes.

4. 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.

If the management network has DHCP…	Then…
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. ⓘ Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

5. Using System Manager on your laptop or console, configure your cluster:

    a. Point your browser to the node management IP address.

    ⓘ The format for the address is https://x.x.x.x.

    b. Configure the system using the data you collected in the ONTAP Configuration Guide

6. Set up your account and download Active IQ Config Advisor:

    a. Log in to your existing account or create an account.

    NetApp Support Registration

    b. Register your system.

    NetApp Product Registration

    c. Download Active IQ Config Advisor.

    NetApp Downloads: Config Advisor

7. Verify the health of your system by running Config Advisor.

8. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain ASA A900 hardware

Maintain the hardware of your ASA A900 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the ASA A900 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the ASA A900 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
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.
DCPM	The DCPM (destage controller power module) contains the NVRAM11 battery.
Fan	The fan cools the controller.
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.
LED USB	The LED USB module provides connectivity to console ports and system status.
NVRAM	The NVRAM module (Non-Volatile Random Access Memory) allows the controller to retain data across power cycles or system reboots, while the NVRAM DIMM maintains NVRAM settings.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - ASA A900

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your ASA A900 storage system.

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

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

1 **Review the boot media requirements**

Review the requirements for boot media replacement.

2 **Shut down the controller**

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

3 **Replace the boot media**

Remove the failed boot media from the controller module and install the replacement boot media.

4 **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5 **Return the failed part to NetApp**

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

Requirements for automated boot media recovery - ASA A900

Before replacing the boot media in your ASA A900, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

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

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.

- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
    - /cfcard/kmip/servers.cfg file.
    - /cfcard/kmip/certs/client.crt file.
    - /cfcard/kmip/certs/client.key file.
    - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
    - The *impaired controller* is the controller on which you are performing maintenance.
    - The *healthy controller* is the HA partner of the impaired controller.

What's next

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

Shut down the controller for automated boot media recovery - ASA A900

Shut down the impaired controller in your ASA A900 storage system to prevent data loss and ensure system stability when replacing the boot media.

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

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

What's next

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

Replace the boot media for automated boot recovery - ASA A900

The boot media in your ASA A900 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module, removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

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

The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

Steps

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

2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.

3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

❶	Cam handle release button
❷	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

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

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

❶	Controller module cover locking button

6. Replace the boot media:

   a. Lift the black air duct at the back of the controller module and then locate the boot media using the following illustration or the FRU map on the controller module:

   [Animation - Replace boot media]()
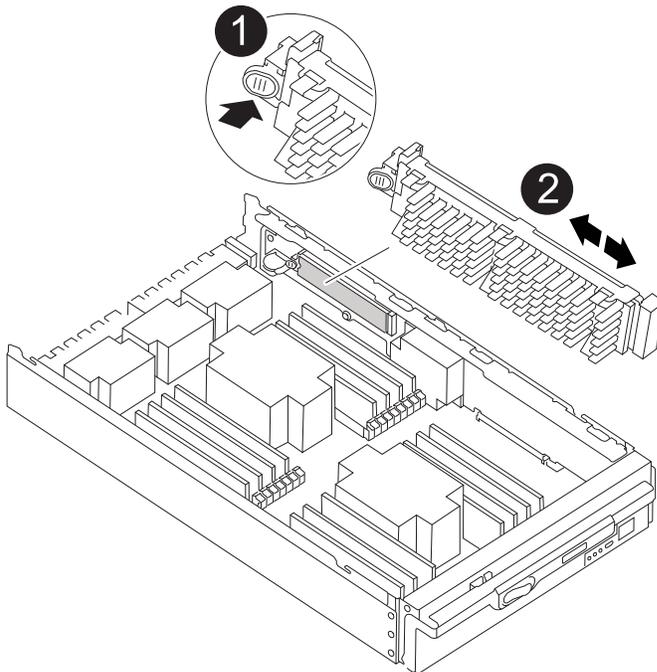
❶	Press release tab
❷	Boot media

   b. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   > ⓘ Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

   c. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.

   d. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

   e. Push the boot media down to engage the locking button on the boot media housing.

7. Reinstall the controller module lid by aligning the pins on the lid with the slots on the motherboard carrier, and then slide the lid into place.

8. Reinstall the controller module:

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

   b. Recable the controller module, as needed.

   c. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, and then push the cam handle to the closed position.

      The controller begins to boot as soon as it is completely installed into the chassis.

      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.

9. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:

   a. Boot to Maintenance mode: `boot_ontap maint`

   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t iniitator adapter_name`

   c. Halt to return to Maintenance mode: `halt`

What's next

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

Automated boot media recovery from the partner node - ASA A900

After installing the new boot media device in your ASA A900 system, you can start the automated boot media recovery process to restore the configuration from the partner node.

During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.
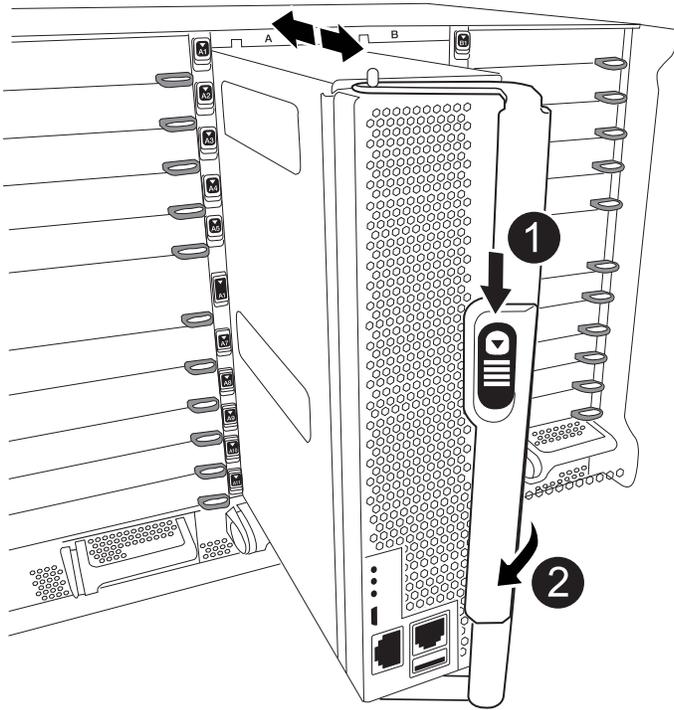
The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

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

Steps

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

```
boot_recovery -partner
```

The screen displays the following message:

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

2. Monitor the boot media install recovery process.

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

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

If you see this message…	Do this…

`key manager is not configured. Exiting.`	Encryption is not installed on the system.
	
	a. Wait for the login prompt to display.
	
	b. Log into the node and give back the storage:
	
	    `storage failover giveback -ofnode`
	    `impaired_node_name`
	
	c. Go to re-enabling automatic giveback if it was disabled.

`key manager is configured.`	Encryption is installed. Go to restoring the key manager.

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

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

Onboard Key Manager (OKM)

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

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

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

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

b. Enter the passphrase for onboard key management when prompted.

c. Enter the passphrase again when prompted to confirm.

d. Enter the backup data for onboard key manager when prompted.

Show example of passphrase and backup data prompts

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

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

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

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

Successfully recovered keymanager secrets.
```

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

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

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

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

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

Go to re-enabling automatic giveback if it was disabled.

External Key Manager (EKM)

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

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

a. Enter the EKM configuration settings when prompted:

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

Show example of client certificate contents

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

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

Show example of client key file contents

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

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

Show example of KMIP server file contents

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

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

Show example of server configuration file contents

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

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

Show example of ONTAP Cluster UUID prompt

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


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

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

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

Show example of temporary network setting prompts

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

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

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

b. Verify the key restoration status:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored. The process restores the appropriate files from the partner node and reboots the node. Proceed to the next step.

- If the key is not successfully restored, the system halts and displays error and warning messages. Rerun the recovery process from the LOADER prompt: `boot_recovery -partner`

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

Uptime: 11m32s
System halting...

LOADER-B>
```

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

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

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

Go to re-enabling automatic giveback if it was disabled.

5. If automatic giveback was disabled, reenable it:

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

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

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

What's next

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

Return the failed part to NetApp - ASA A900

If a component in your ASA A900 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Replace the boot media - ASA A900

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` file.

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

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair does not require connection to a network to restore the `var` file system. The HA pair in a single chassis has an internal e0S connection, which is used to transfer `var` config between them.
  - For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from NetApp.
- 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 encryption key support and status - ASA A900

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

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

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

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

Steps

1. Check if your ONTAP version supports encryption:

   ```
   version -v
   ```

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Download the appropriate ONTAP image based on NVE support:
   - If NVE is supported: Download the ONTAP image with NetApp Volume Encryption

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

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

Step 2: Verify key manager status and back up configuration

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

Steps

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

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

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

   If no key manager is configured:

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

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

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

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

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

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

External Key Manager (EKM)

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

If all keys show `true` in the Restored column:

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

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

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.

b. Verify that all authentication keys are restored:

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

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

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

Onboard Key Manager (OKM)

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

If all keys show `true` in the Restored column:

a. Back up the OKM information:

    i. Switch to advanced privilege mode:

```
set -priv advanced
```

    Enter `y` when prompted to continue.

    ii. Display the key management backup information:

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

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

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

    iv. Return to admin mode:

```
set -priv admin
```

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

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

a. Synchronize the onboard key manager:

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

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

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

b. Verify all authentication keys are restored:

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

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

c. Back up the OKM information:

    i. Switch to advanced privilege mode:

```
set -priv advanced
```

    Enter `y` when prompted to continue.

    ii. Display the key management backup information:

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

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

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

    iv. Return to admin mode:

```
set -priv admin
```

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

Shut down the controller for manual boot media recovery - ASA A900

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

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

Most configurations

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Controller is in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

ⓘ 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 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 `y` when prompted.

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

Replace the boot media and prepare for manual boot recovery - ASA A900

You must unplug the controller module, remove and open the controller module, locate and replace the boot media in the controller, and then transfer the image to the replacement boot media.

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

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

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

2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.

3. Slide the terra cotta button on the cam handle downward until it unlocks.

Animation - Remove the controller

❶	Cam handle release button
❷	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

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

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

①	Controller module cover locking button

Step 2: Replace the boot media

You must locate the boot media in the controller and follow the directions to replace it.

Steps

1. Lift the black air duct at the back of the controller module and then locate the boot media using the following illustration or the FRU map on the controller module:

   [Animation - Replace boot media](#)
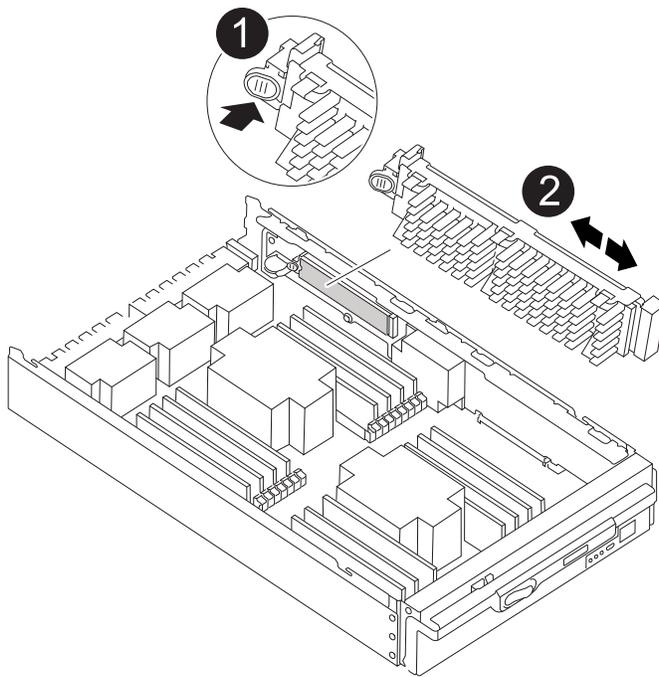
①	Press release tab
②	Boot media

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   > ⓘ Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.

4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

6. Reinstall the controller module lid by aligning the pins on the lid with the slots on the motherboard carrier, and then slide the lid into place.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the var file system during this procedure.

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.

- Download a copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site. Use the `version -v` command to display if your version of ONTAP supports NVE. If the command output displays `<10no- DARE>`, your version of ONTAP does not support NVE.

  ◦ If NVE is supported by your version of ONTAP, download the image with NetApp Volume Encryption, as indicated in the download button.

  ◦ If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. If you have not done so, 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.

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.

3. Recable the controller module, as needed.

4. Insert the USB flash drive into the USB slot on the controller module.

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

5. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, and then push the cam handle to the closed position.

   The controller begins to boot as soon as it is completely installed into the chassis.

6. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

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.

7. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:

   a. Boot to Maintenance mode: `boot_ontap maint`

   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t initiator adapter_name`

   c. Halt to return to Maintenance mode: `halt`

   The changes will be implemented when the system is booted.

Manual boot media recovery from a USB drive - ASA A900

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

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

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

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

   `boot_recovery`

   The recovery image is downloaded from the USB flash drive.

2. When prompted, enter the name of the image or press **Enter** to accept the default image displayed in brackets.

3. Restore the var file system using the procedure for your ONTAP version:

ONTAP 9.16.0 or earlier

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

a. **On the impaired controller:** Press `Y` when you see `Do you want to restore the backup configuration now?`

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

c. **On the partner controller:** Set the impaired controller to advanced privilege level:

```
set -privilege advanced
```

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

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

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

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

```
set -privilege admin
```

f. **On the impaired controller:** Press `Y` when you see `Was the restore backup procedure successful?`

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

h. **On the impaired controller:** Press `Y` when prompted to reboot, then press `Ctrl-C` when you see the Boot Menu.

i. **On the impaired controller:** Do one of the following:

   - If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.
   - If the system uses encryption, go to Restore encryption.

ONTAP 9.16.1 or later

Complete the following steps on the impaired controller:

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

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

b. Press `Y` when prompted to confirm that the restore backup was successful.

c. Press `Y` when prompted to use the restored configuration.

d. Press `Y` when prompted to reboot the node.

e. Press `Y` when prompted to reboot again, then press `Ctrl-C` when you see the Boot Menu.

f. Do one of the following:

   - If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.

- If the system uses encryption, go to Restore encryption.

4. Connect the console cable to the partner controller.

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

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

6. If you disabled automatic giveback, reenable it:

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

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

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

Restore encryption - ASA A900

Restore encryption on the replacement boot media.

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

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

Onboard Key Manager (OKM)

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

Before you begin

Ensure you have the following information available:

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

Steps

On the impaired controller:

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

ONTAP version	Select this option

ONTAP 9.8 or later	Select option 10.
	
	Show example boot menu
	
	<pre>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</pre>

ONTAP version	Select this option
ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager` **Show example boot menu** ``` 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. Selection (1-19)? recover_onboard_keymanager ```

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

   Show example prompt

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

4. Enter the cluster-wide passphrase twice.

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

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information:

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

```
Enter the backup data:

---------------------------BEGIN
BACKUP-------------------------
01234567890123456789012345678901234567890123456789012345678901
23
12345678901234567890123456789012345678901234567890123456789012
34
23456789012345678901234567890123456789012345678901234567890123
45
34567890123456789012345678901234567890123456789012345678901234
56
45678901234567890123456789012345678901234567890123456789012345
67
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
```

```
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
01234567890123456789012345678901234567890123456789012345678901
23
12345678901234567890123456789012345678901234567890123456789012
34
23456789012345678901234567890123456789012345678901234567890123
45
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA

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

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

The recovery process completes and displays the following message:

```
Successfully recovered keymanager secrets.
```

Show example prompt

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

Successfully recovered keymanager secrets.


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

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

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

Show example prompt

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



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

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

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

   On the partner controller:

8. Giveback the impaired controller:

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

   On the impaired controller:

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

   `security key-manager onboard sync`

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

Show example prompt

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

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

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

11. Verify that all keys are synced:

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

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

    On the partner controller:

12. Giveback the impaired controller:

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

13. Restore automatic giveback if you disabled it:

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

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

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

External Key Manager (EKM)

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

Before you begin

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

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

Steps

On the impaired controller:

1. Connect the console cable to the impaired controller.

2. Select option `11` from the ONTAP boot menu.

   Show example boot menu

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

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

   Show example prompt

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

4. Enter the client and server information when prompted:

   a. Enter the client certificate (client.crt) file contents, including the BEGIN and END lines.

   b. Enter the client key (client.key) file contents, including the BEGIN and END lines.

   c. Enter the KMIP server CA(s) (CA.pem) file contents, including the BEGIN and END lines.

   d. Enter the KMIP server IP address.

   e. Enter the KMIP server port (press Enter to use the default port 5696).

Show example

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

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

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

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

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

The recovery process completes and displays the following message:

```
Successfully recovered keymanager secrets.
```

Show example

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

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

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

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

6. Restore automatic giveback if you disabled it:

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

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

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

Return the failed part to NetApp - ASA A900

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

Chassis

Replace the chassis - ASA A900

To replace the chassis, you must remove the power supplies, fans, controller modules, I/O modules, DCPM modules, and USB LED module from the impaired chassis, remove the impaired chassis from the equipment rack or system cabinet, install the replacement chassis in its place, and then install the components into the replacement chassis.

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 disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shutdown the controllers - ASA A900

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

Before you begin

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

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.

2. Stop all clients/host from accessing data on the NetApp system.

3. Suspend external backup jobs.

4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

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

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

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

6. Exit the cluster shell:

   ```
   exit
   ```

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

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

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

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

> (i) For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict -sync-warnings true`

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

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

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

Move and replace hardware - ASA A900

To replace the chassis, you must remove the components from the impaired chassis and install them in the replacement chassis.

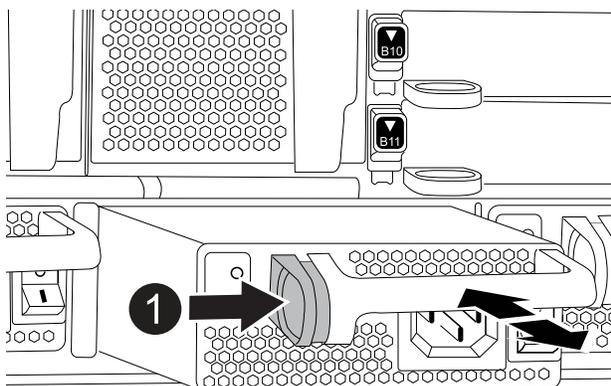**Step 1: Remove the power supplies**

Removing the power supplies when replacing a chassis involves turning off, disconnecting, and then removing the four power supplies from the rear of the impaired chassis.

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

2. Turn off the power supply and disconnect the power cables:

   a. Turn off the power switch on the power supply.

   b. Open the power cable retainer, and then unplug the power cable from the power supply.

   c. Unplug the power cable from the power source.

3. Press and hold the terra cotta locking button on the power supply handle, and then pull the power supply out of the chassis.

   > (i) When removing a power supply, always use two hands to support its weight.

Animation - Remove/install PSU

①	Locking button

4. Repeat the preceding steps for any remaining power supplies.
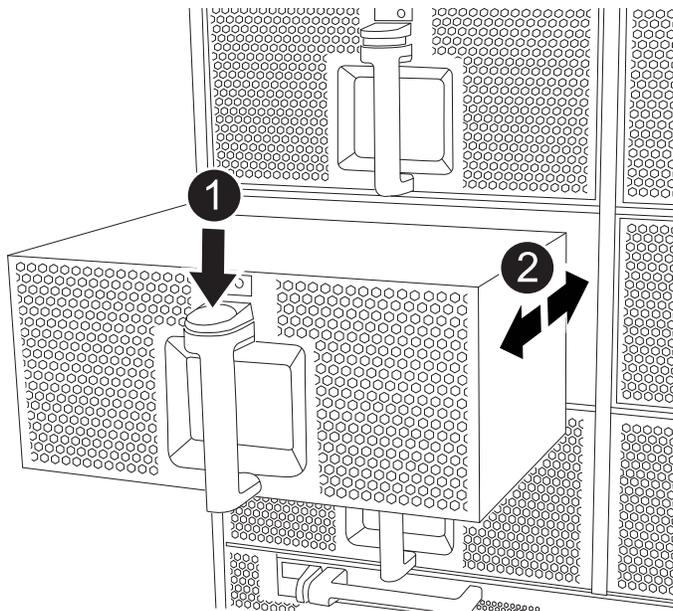
Step 2: Remove the fans

You must remove the six fan modules, located on in the front of the chassis, when replacing the chassis.

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

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

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

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

Animation - Remove/install fan

| ①	Terra cotta locking button
②	Slide fan in/out of chassis

4. Set the fan module aside.

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

Step 3: Remove the controller module

To replace the chassis, you must remove the controller module or modules from the impaired chassis.

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta locking button on the cam handle downward until it unlocks.

Animation - Remove the controller
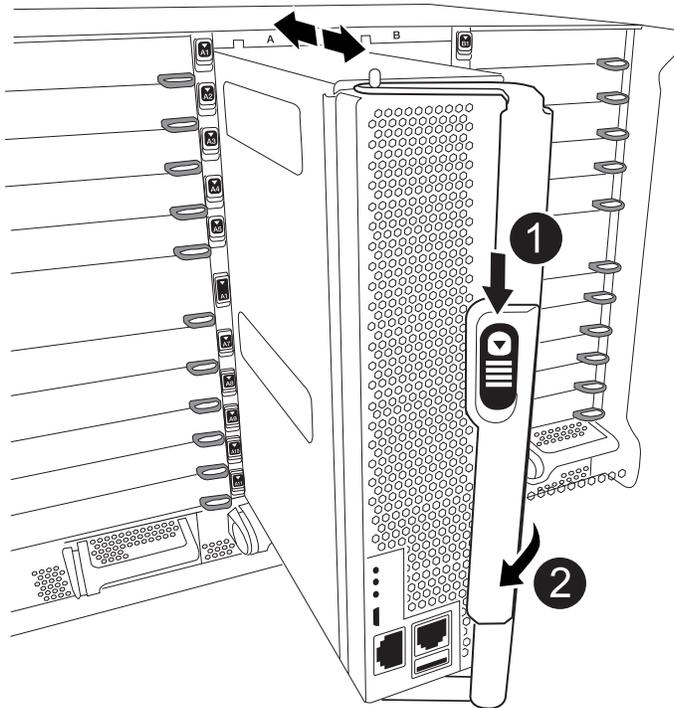
1	Cam handle locking button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

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

5. Set the controller module aside in a safe place and keep track of which chassis slot it came from, so that it can be installed into the same slot in the replacement chassis..

6. Repeat these steps if you have another controller module in the chassis.

Step 4: Remove the I/O modules

To remove I/O modules from the impaired chassis, including the NVRAM modules, follow the specific sequence of steps.

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

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

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

   a. Depress the lettered and numbered cam locking button.

   The cam locking button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

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

   Animation - Remove/install I/O module

|	
❶	Lettered and numbered I/O cam latch
❷	I/O cam latch completely unlocked
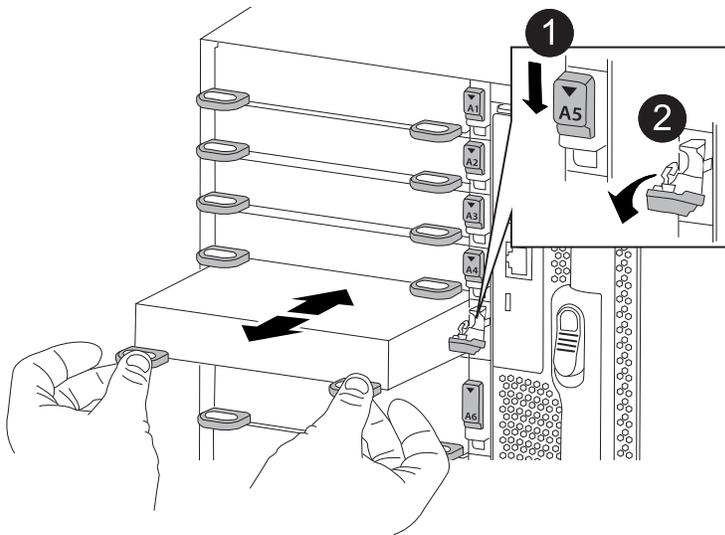
4. Set the I/O module aside.
5. Repeat the preceding step for the remaining I/O modules in the impaired chassis.

Step 5: Remove the de-stage controller power module

Remove the two de-stage controller power modules from the front of the impaired chassis.

1. If you are not already grounded, properly ground yourself.
2. Press the terra cotta locking button on the module handle, and then slide the DCPM out of the chassis.
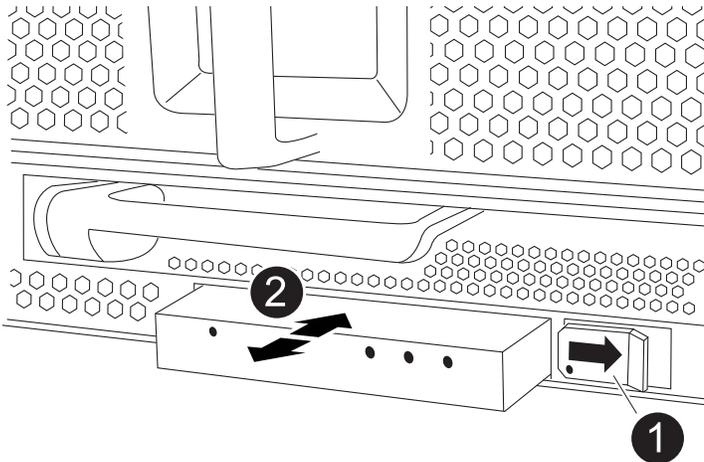
   Animation - Remove/install DCPM
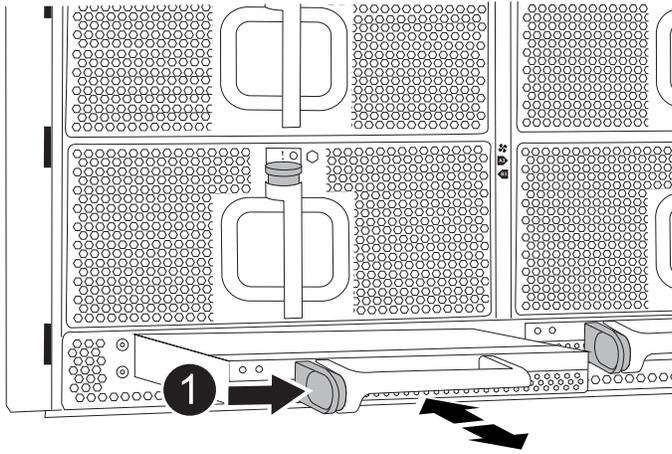
1	DCPM terra cotta locking button

3. Set the DCPM aside in a safe place and repeat this step for the remaining DCPM.

Step 6: Remove the USB LED module

Remove the USB LED modules.

[Animation - Remove/install USB](#)

1	Eject the module.
2	Slide out of chassis.

1. Locate the USB LED module on the front of the impaired chassis, directly under the DCPM bays.

2. Press the black locking button on the right side of the module to release the module from the chassis, and then slide it out of the impaired chassis.

3. Set the module aside in a safe place.

Step 7: Remove chassis

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.

   > (i) If the system is in a system cabinet, you might need to remove the rear tie-down bracket.

2. With the help of two or three people, slide the impaired chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.

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

4. Using two or three 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 *L* brackets in an equipment rack.

5. Slide the chassis all the way into the equipment rack or system cabinet.

6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the impaired chassis.

7. Secure the rear of the chassis to the equipment rack or system cabinet.

8. If you are using the cable management brackets, remove them from the impaired chassis, and then install them on the replacement chassis.

Step 8: Install the de-stage controller power module

When the replacement chassis is installed into the rack or system cabinet, you must reinstall the de-stage controller power modules into it.

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

2. Align the end of the DCPM with the chassis opening, and then gently slide it into the chassis until it clicks into place.

   > (i) The module and slot are keyed. Do not force the module into the opening. If the module does not go in easily, realign the module and slide it into the chassis.

3. Repeat this step for the remaining DCPM.

Step 9: Install fans into the chassis

To install the fan modules when replacing the chassis, you must perform a specific sequence of tasks.

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

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

   When inserted into a live system, the amber Attention LED flashes four times when the fan module is successfully inserted into the chassis.

3. Repeat these steps for the remaining fan modules.

4. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.

Step 10: Install I/O modules

To install I/O modules, including the NVRAM modules from the impaired chassis, follow the specific sequence of steps.

You must have the chassis installed so that you can install the I/O modules into the corresponding slots in the replacement chassis.

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

2. After the replacement chassis is installed in the rack or cabinet, install the I/O modules into their corresponding slots in the replacement chassis by gently sliding the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage, and then push the I/O cam latch all the way up to lock the module in place.

3. Recable the I/O module, as needed.

4. Repeat the preceding step for the remaining I/O modules that you set aside.

    (i) If the impaired chassis has blank I/O panels, move them to the replacement chassis at this time.

Step 11: Install the power supplies

Installing the power supplies when replacing a chassis involves installing the power supplies into the replacement chassis, and connecting to the power source.

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

2. Make sure the power supplies rockers are in the off position.

3. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis until it locks into place.

    The power supplies are keyed and can only be installed one way.

    (i) Do not use excessive force when sliding the power supply into the system. You can damage the connector.

4. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

    (i) Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

5. Repeat the preceding steps for any remaining power supplies.

Step 12: Install the USB LED modules

Install the USB LED modules in the replacement chassis.

1. Locate the USB LED module slot on the front of the replacement chassis, directly under the DCPM bays.

2. Align the edges of the module with the USB LED bay, and gently push the module all the way into the chassis until it clicks into place.

Step 13: Install the controller

After you install the controller module and any other components into the replacement chassis, boot it.

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

2. Connect the power supplies to different power sources, and then turn them on.

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

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

4. Recable the console to the controller module, and then reconnect the management port.

5. With the cam handle in the open position, slide the controller module into the chassis and firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle until it clicks into the locked position.

   > ℹ️   Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors.

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

6. Repeat the preceding steps to install the second controller into the replacement chassis.

7. Boot each controller.

Restore and verify the configuration - ASA A900

To complete the chassis replacement, you must complete specific tasks.

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
      - non-ha

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

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

Step 2: Bring up the system

1. If you have not done so, plug the power cables back into the PSUs.

2. Turn on the PSUs by toggling the rocker switched to **ON**, and wait for the controllers to power up completely.

3. Check the front and the back of the chassis and controllers for any fault lights after power up.

4. Connect to the SP or BMC IP address of the nodes via SSH. This will be the same address used to shut down the nodes.

5. Perform additional health checks as described in How_to_perform_a_cluster_health_check_with_a_script_in_ONTAP

6. Turn AutoSupport back on (end the maintenance window message): `system node autosupport invoke -node * -type all -message MAINT=end`

> (i) As a best practice, you should do the following:
>
> • Resolve any Active IQ Wellness Alerts and Risks (Active IQ will take time to process post-power up AutoSupports - expect a delay in results)
>
> • Run Active IQ Config Advisor
>
> • Check system health using How_to_perform_a_cluster_health_check_with_a_script_in_ONTAP

Step 3: Return the failed part to NetApp

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

Controller

Replace the controller module - ASA A900

To replace the impaired controller module, you must shut down the impaired controller, move the internal components to the replacement controller module, install the replacement controller module, and reboot the replacement controller.

Before you begin
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.

- If your system has a V_StorageAttach license, you must refer to the additional required steps before performing this procedure.

- If your system is in an HA pair, 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.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific

steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- 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.

- In this procedure, the boot device is moved from the impaired controller to the replacement controller so that the replacement controller will boot up in the same version of ONTAP as the old 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 file.

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

Shut down the impaired controller - ASA A900

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

If the impaired controller is displaying…	Then…

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

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 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 `y` when prompted.

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

Replace the controller module hardware - ASA A900

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

The following animation shows the whole process of moving components from the impaired to the replacement controller.
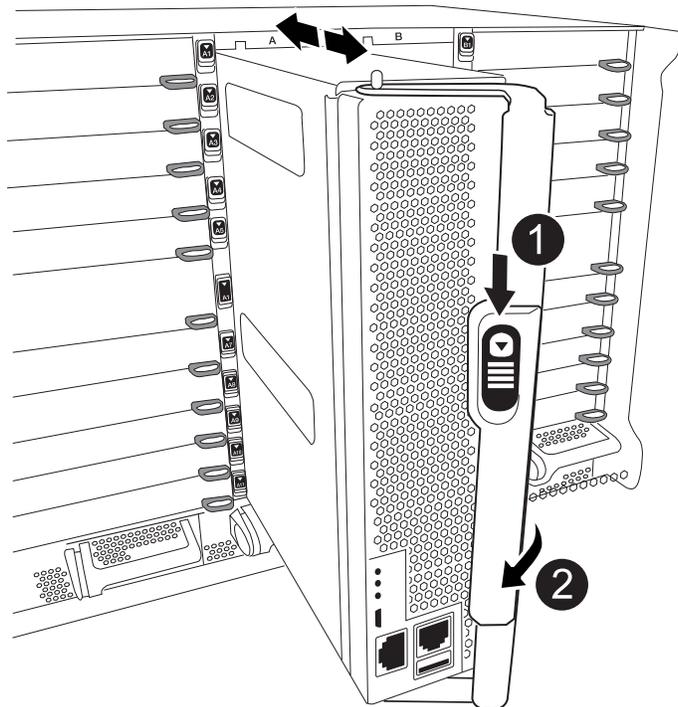
Animation - Move components to replacement controller

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

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

2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.

3. Slide the terra cotta button on the cam handle downward until it unlocks.
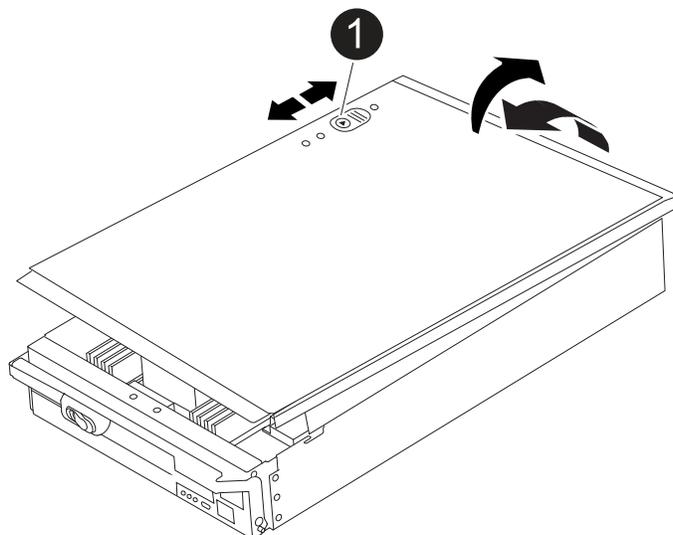
   Animation - Remove the controller

❶	Cam handle release button
❷	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

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

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.
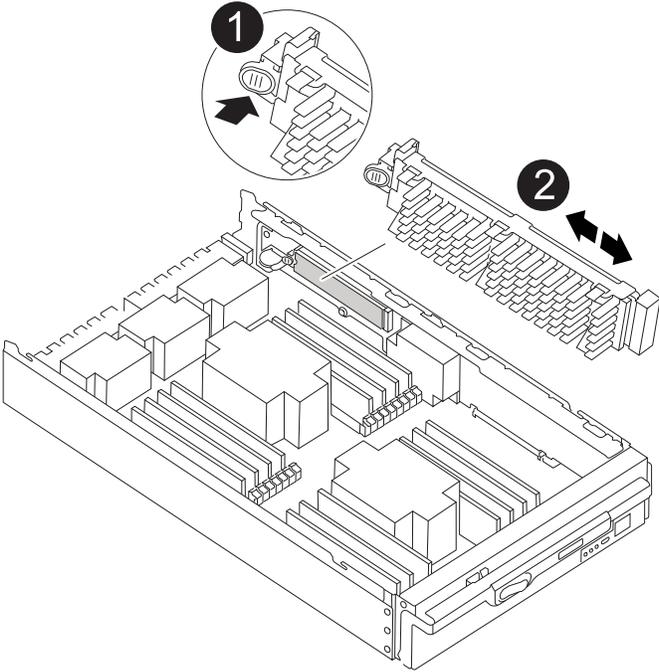
①	Controller module cover locking button

Step 2: Move the boot media

You must locate the boot media and follow the directions to remove it from the old controller and insert it in the new controller.

1. Locate the boot media using the following illustration or the FRU map on the controller module:

①	Press release tab
②	Boot media

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   ⓘ Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.

4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

Step 3: Move the system DIMMs

To move the DIMMs, locate and move them from the old controller into the replacement controller and follow the specific sequence of steps.

> (i) The VER2 controller has fewer DIMM sockets. There is no reduction in the number of DIMMs supported or change in the DIMM socket numbering. When moving the DIMMs to the new controller module, install the DIMMs into the same socket number/location as the impaired controller module. See the FRU map diagram on the VER2 controller module for DIMM socket locations.

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

2. Locate the DIMMs on your controller module.

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.

   > (i) Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.
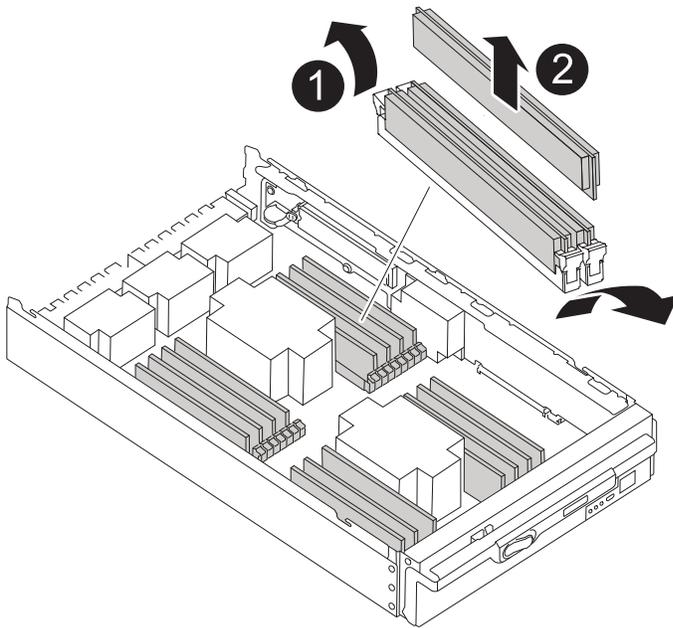
1	DIMM ejector tabs
2	DIMM

5. Locate the slot where you are installing the DIMM.

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

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

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

7. 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. |

8. 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.

9. Repeat these steps for the remaining DIMMs.

Step 4: Install the controller

After you install the components into the replacement controller module, you must install the replacement controller module into the system chassis and boot the operating system.
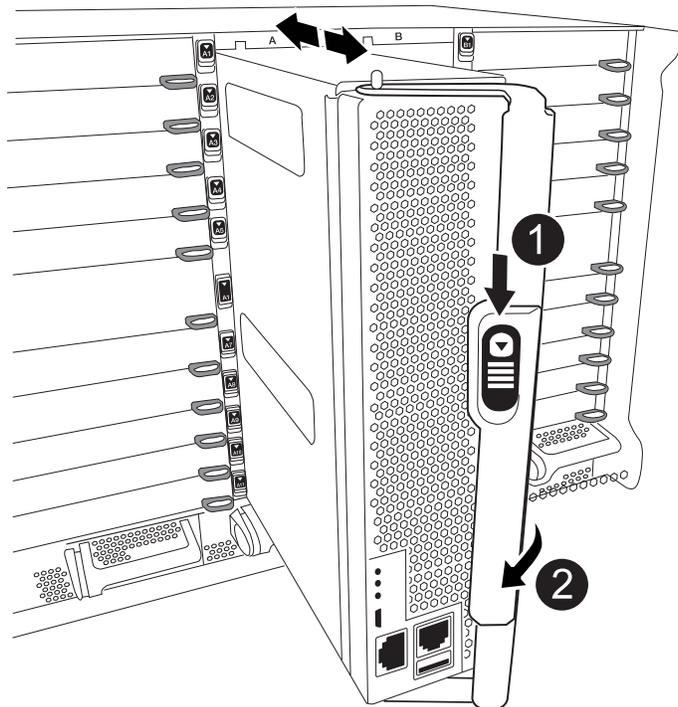
For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

| ⓘ | The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process. |

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

2. If you have not already done so, replace the cover on the controller module.

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

Animation - Install controller

1	Cam handle release button
2	Cam handle

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

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

> ℹ️ You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:

   a. If you have not already done so, reinstall the cable management device.

   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

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

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   c. Rotate the controller module cam handle to the locked position.

   d. Interrupt the boot process by pressing `Ctrl-C` when you see Press Ctrl-C for Boot Menu.

   e. Select the option to boot to LOADER.

Restore and verify the system configuration - ASA A900

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 the 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 controller's HA state

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 replacement controller module, verify that all components display the same `HA` state: `ha-config show`

If your system is in…	The HA state for all components should be…
An HA pair	ha
A MetroCluster FC configuration with four or more nodes	mcc

If your system is in…	The HA state for all components should be…
A MetroCluster IP configuration	mccip

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`

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

Recable the system - ASA A900

Continue the replacement procedure by recabling the storage and network coniguration.

Step 1: Recable the system

You must recable the controller module's storage and network connections.

Steps

1. Recable the system.

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.

   > ⓘ The system ID and disk assignment information reside in the NVRAM module, which is in a module separate from the controller module and not impacted by the controller module replacement.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* 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 *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`

3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

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

replacement and has a new system ID of 151759706.

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

4. From the healthy controller, verify that any coredumps are saved:

   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node` *local-node-name* `partner savecore`

   c. Wait for the savecore command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node` *local-node-name* `partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`

5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:

   ◦ Restore onboard key management encryption keys
   ◦ Restore external key management encryption keys

6. Give back the controller:

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

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

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

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

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

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

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

   ```
   node1> `storage disk show -ownership`

   Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID   DR Home ID
   Reserver  Pool
   ----- ------    ----- ------ -------- -------    -------    -------
   --------- ---
   1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277 -
   1873775277 Pool0
   1.0.1  aggr0_1  node1 node1           1873775277 1873775277 -
   1873775277 Pool0
   .
   .
   .
   ```

8. 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.

9. 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.

     For more information, see Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration topic.

10. 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.
```

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

12. 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 - ASA A900

To complete the replacement procedure and restore your system to full operation, you must recable the storage, restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller. You must complete a series of tasks before restoring your system to full operation.

Step 1: Install licenses for the new controller

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

Before you begin

If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on ONTAP platforms. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

About this task

- Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

  Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

- The 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.

- If the node is in a MetroCluster configuration and all nodes at a site have been replaced, license keys must be installed on the *replacement* node or nodes prior to switchback.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support site in the My Support section under Software licenses.

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

2. Install each license key: `system license add -license-code license-key, license-key...`

3. Remove the old licenses, if desired:

   a. Check for unused licenses: `license clean-up -unused -simulate`

   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and register the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

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

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

2. Register the system serial number with NetApp Support.

   ◦ If AutoSupport is enabled, send an AutoSupport message to register the serial number.

   ◦ If AutoSupport is not enabled, call NetApp Support to register the serial number.

3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.

4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

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

Step 3: Return the failed part to NetApp

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

Replace a DIMM - ASA A900

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

Before you begin

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:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

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 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 `y` when prompted.

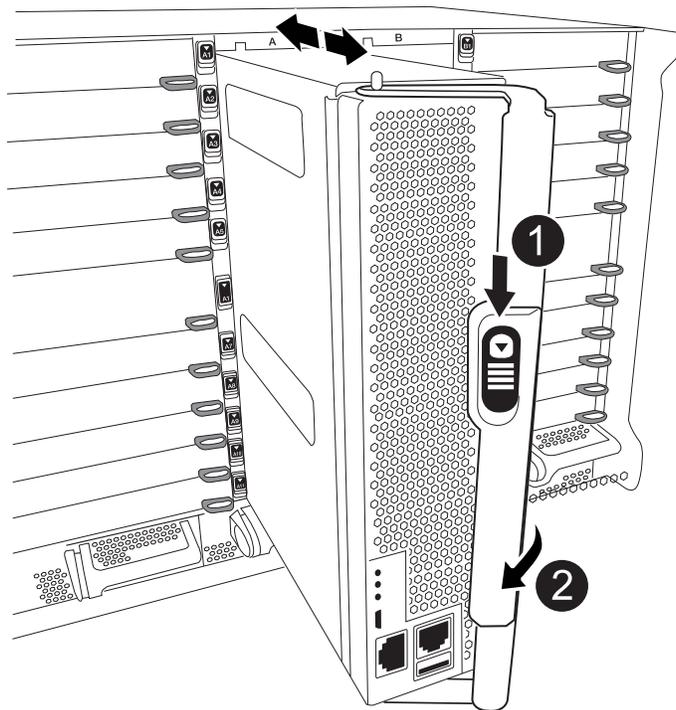| If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press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

Step 2: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

Animation - Remove the controller

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

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

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

1	Controller module cover locking button

Step 3: Replace the DIMMs

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

> ⓘ The VER2 controller has fewer DIMM sockets. There is no reduction in the number of DIMMs supported or change in the DIMM socket numbering. When moving the DIMMs to the new controller module, install the DIMMs into the same socket number/location as the impaired controller module. See the FRU map diagram on the VER2 controller module for DIMM socket locations.

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller module.

1   2   3   4   5   6   7   8   9   10   11   12   13   14   15   16

17   18   19   20   21   22   23   24   25   26   27   28   29   30   31   32
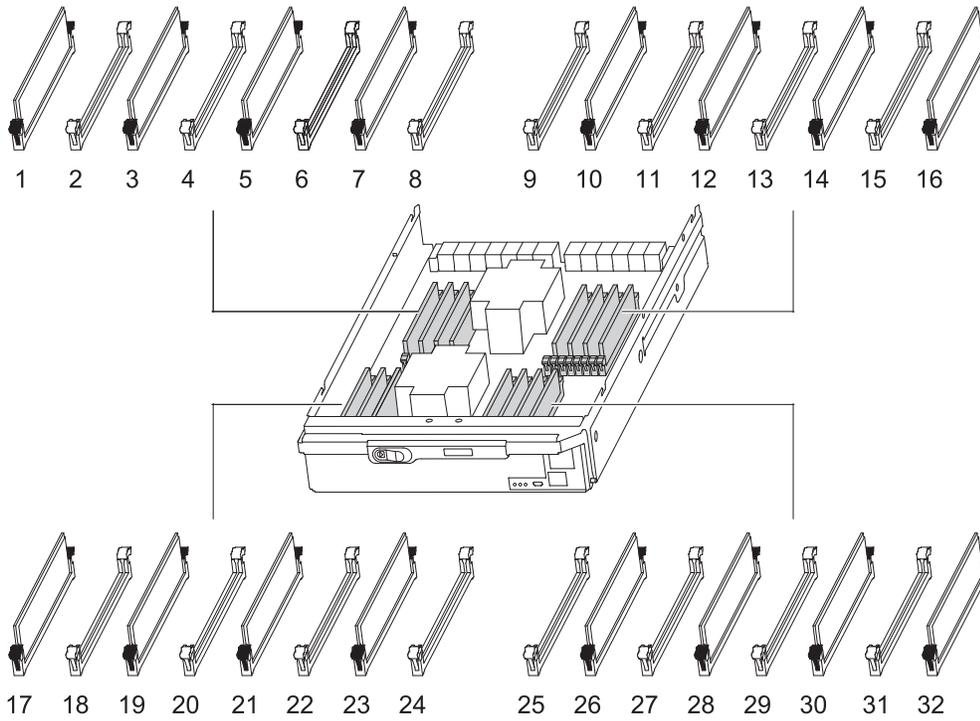
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

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

Animation - Replace DIMM
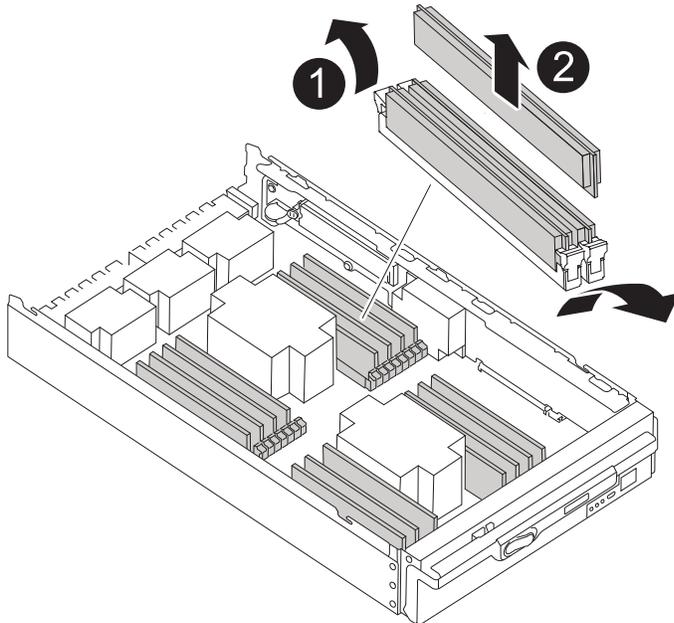
❶	DIMM ejector tabs

②	DIMM

4. 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.

5. 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.

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

7. Close the controller module cover.

Step 4: Install the controller

After you install the components into the controller module, you must install the controller module back into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

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

2. If you have not already done so, replace the cover on the controller module.

❶	Controller module cover locking button

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

❶	Cam handle release button
❷	Cam handle

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

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

> ℹ️ You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:

   a. If you have not already done so, reinstall the cable management device.

   b. 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.

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

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them

into the locked position.

    d. Interrupt the boot process by pressing `Ctrl-C` when you see Press Ctrl-C for Boot Menu.

    e. Select the option to boot to Maintenance mode from the displayed menu.

Step 5: Run system-level diagnostics

After installing a new DIMM, you should run diagnostics.

Your system must be at the LOADER prompt to start System Level Diagnostics.

All commands in the diagnostic procedures are issued from the controller where the component is being replaced.

1. If the controller to be serviced is not at the LOADER prompt, perform the following steps:

    a. Select the Maintenance mode option from the displayed menu.

    b. After the controller boots to Maintenance mode, halt the controller: `halt`

    After you issue the command, you should wait until the system stops at the LOADER prompt.

> ⓘ   During the boot process, you can safely respond `y` to prompts.

      ▪ If a prompt appears warning that when entering Maintenance mode in an HA configuration, you must ensure that the healthy controller remains down.

2. At the LOADER prompt, access the special drivers specifically designed for system-level diagnostics to function properly: `boot_diags`

   During the boot process, you can safely respond `y` to the prompts until the Maintenance mode prompt (*>) appears.

3. Run diagnostics on the system memory: `sldiag device run -dev mem`

4. Verify that no hardware problems resulted from the replacement of the DIMMs: `sldiag device status -dev mem -long -state failed`

   System-level diagnostics returns you to the prompt if there are no test failures, or lists the full status of failures resulting from testing the component.

5. Proceed based on the result of the preceding step:

If the system-level diagnostics tests…	Then…

Were completed without any failures	a. Clear the status logs: `sldiag device clearstatus`
	
	b. Verify that the log was cleared: `sldiag device status`
	
	   The following default response is displayed:
	
	   SLDIAG: No log messages are present.
	
	c. Exit Maintenance mode: `halt`
	
	   The controller displays the LOADER prompt.
	
	d. Boot the controller from the LOADER prompt: `bye`
	
	e. Return the controller to normal operation:

If your controller is in…	Then…
An HA pair	Perform a give back: `storage failover giveback -ofnode replacement_node_name` **Note:** If you disabled automatic giveback, re-enable it with the storage failover modify command.

If your controller is in…	Then…

Resulted in some test failures	Determine the cause of the problem:
	
	a. Exit Maintenance mode: `halt`
	
	   After you issue the command, wait until the system stops at the LOADER prompt.
	
	b. Verify that you have observed all the considerations identified for running system-level diagnostics, that cables are securely connected, and that hardware components are properly installed in the storage system.
	
	c. Boot the controller module you are servicing, interrupting the boot by pressing `Ctrl-C` when prompted to get to the Boot menu:
	
	   ◦ If you have two controller modules in the chassis, fully seat the controller module you are servicing in the chassis.
	
	     The controller module boots up when fully seated.
	
	   ◦ If you have one controller module in the chassis, connect the power supplies, and then turn them on.
	
	d. Select Boot to maintenance mode from the menu.
	
	e. Exit Maintenance mode by entering the following command: `halt`
	
	   After you issue the command, wait until the system stops at the LOADER prompt.
	
	f. Rerun the system-level diagnostic test.

Step 6: Return the failed part to NetApp

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

Replace the DCPM containing the NVRAM11 battery - ASA A900

To hot-swap a destage controller power module (DCPM), which contains the NVRAM11 battery, you must locate the failed DCPM module, remove it from the chassis, and install the replacement DCPM module.

You must have a replacement DCPM module in-hand before removing the failed module from the chassis and it must be replaced within five minutes of removal. Once the DCPM module is removed from the chassis, there is no shutdown protection for the controller module that owns the DCPM module, other than failover to the other controller module.

Step 1: Replace the DCPM module

To replace the DCPM module in your system, you must remove the failed DCPM module from the system and then replace it with a new DCPM module.

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

2. Remove the bezel on the front of the system and set it aside.

3. Locate the failed DCPM module in the front of the system by looking for the Attention LED on the module.

   The LED will be steady amber if the module is faulty.

   (i) The DCPM module must be replaced in the chassis within five minutes of removal or the associated controller will shut down.

4. Press the terra cotta release button on the module handle, and then slide the DCPM module out of the chassis.

   Animation - Remove/install DCPM
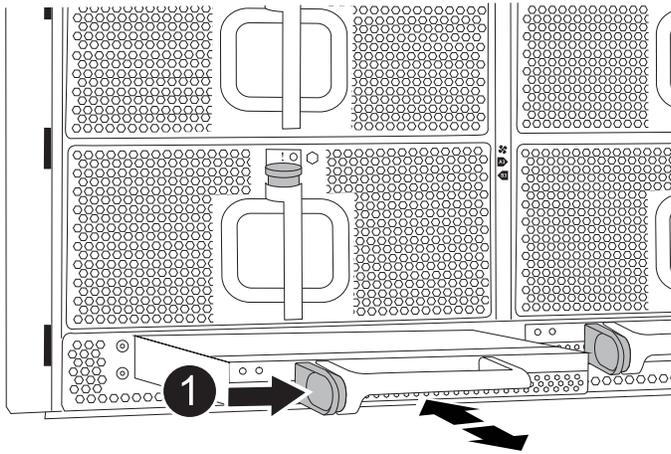
1	DCPM module terra cotta locking button

5. Align the end of the DCPM module with the chassis opening, and then gently slide it into the chassis until it clicks into place.

   (i) The module and slot are keyed. Do not force the module into the opening. If the module does not go in easily, realign the module and slide it into the chassis.

   The Amber LED flashes four times upon insertion and the green LED also flashes if the battery is providing a voltage. If it does not flash, it will likely need to be replaced.

Step 2: Dispose of batteries

You must dispose of batteries according to the local regulations regarding battery recycling or disposal. If you cannot properly dispose of batteries, you must return the batteries to NetApp, as described in the RMA instructions that are shipped with the kit.

Safety Information and Regulatory Notices

Step 3: Return the failed part to NetApp

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

Swap out a fan - ASA A900

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

> ⓘ It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

Steps

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

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

Animation - Remove/install fan
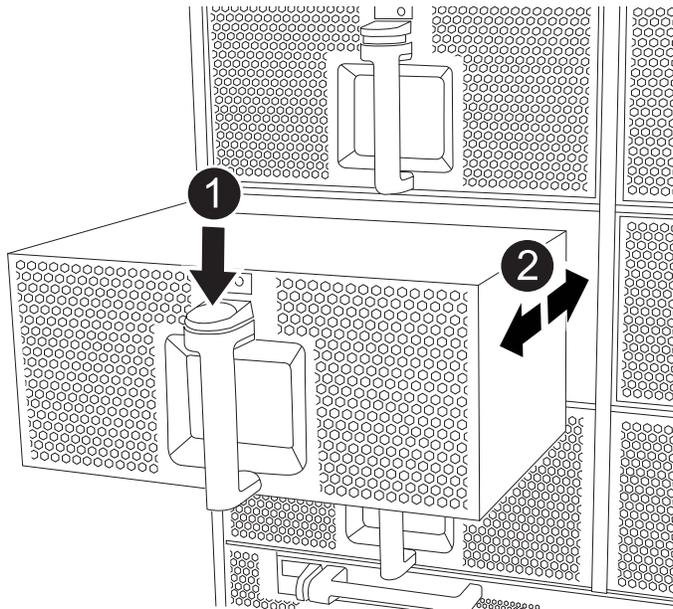
❶	Terra cotta release button

②	Slide fan in/out of chassis

5. Set the fan module aside.

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

   When inserted into a live system, the amber Attention LED flashes four times when the fan module is successfully inserted into the chassis.

7. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.

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

I/O module

Replace the I/O module - ASA A900

To replace an I/O module, you must perform a specific sequence of tasks.

- 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 node

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

Option 1: Most systems

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

Before you begin

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

Steps

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

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

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

   `storage failover modify -node impaired-node -auto-giveback-of false`

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

   a. If the impaired controller cannot be brought up or is already taken over, you must take the HA interconnect link down from the healthy controller before booting up the impaired controller. This will prevent the impaired controller from performing automatic giveback.

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

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

3. Take the impaired controller to the LOADER prompt:

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

Option 2: Controller is in a MetroCluster

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

Step 2: Replace I/O modules

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

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

2. Unplug any cabling associated with the target I/O module.

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

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

a. Depress the lettered and numbered cam button.

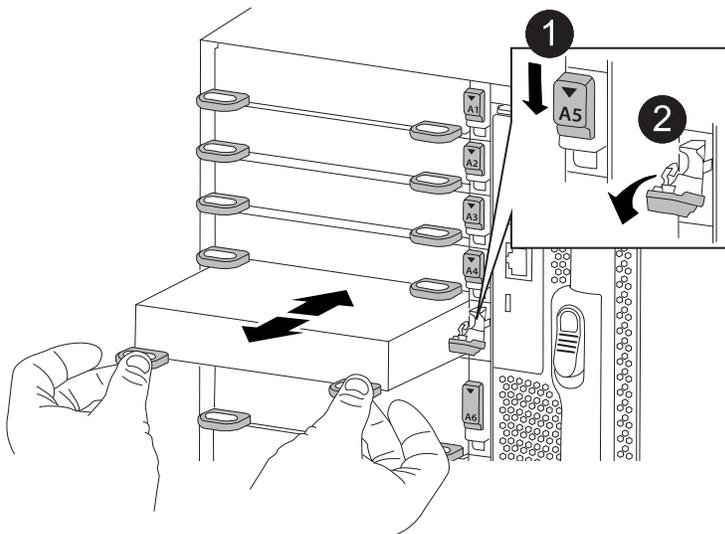   The cam button moves away from the chassis.

b. Rotate the cam latch down until it is in a horizontal position.

   The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

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

   Animation - Remove/install I/O module

①	Lettered and numbered I/O cam latch
②	I/O cam latch completely unlocked

4. Set the I/O module aside.
5. Install the replacement I/O module into the chassis by gently sliding the I/O module into the slot until the lettered and numbered 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.
6. Recable the I/O module, as needed.

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: `priv set advanced`

   b. Reboot the BMC: `sp reboot`

2. From the LOADER prompt, reboot the node: `bye`

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

3. If your system is configured to support 10 GbE cluster interconnect and data connections on 40 GbE NICs, convert these ports to 10 GbE connections by using the nicadmin convert command from Maintenance mode. See Convert 40GbE NIC ports into multiple 10GbE ports for 10GbE connectivity for more information.

   > ℹ️ Be sure to exit Maintenance mode after completing the conversion.

4. Return the node to normal operation: `storage failover giveback -ofnode impaired_node_name`

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

Step 4: Return the failed part to NetApp

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

Add an I/O module - ASA A900

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

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your 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.

- To non-disruptively add an I/O module, you must take over the target controller, remove the slot blanking cover in the target slot or remove an existing I/O module, add the new or replacement I/O module, and then giveback the target controller.

- Make sure that all other components are functioning properly.

Add I/O module to an empty slot

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

Steps

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

2. Remove the target slot blanking cover:

   a. Depress the lettered and numbered cam latch.

   b. Rotate the cam latch down until it is the open position.

   c. Remove the blanking cover.

3. Install the I/O module:

   a. Align the I/O module with the edges of the slot.

   b. Slide the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.

   c. Push the I/O cam latch all the way up to lock the module in place.

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

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

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

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

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

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

8. If you are using slots 3 and/or 7 for networking, use the `storage port modify -node <node name> -port <port name> -mode network` command to convert the slot for networking use.

9. Repeat these steps for controller B.

10. If you installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.

Add I/O module to a fully-populated system

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

About this task

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

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

Steps

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

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

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

   a. Depress the lettered and numbered cam latch.

      The cam latch moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

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

      Animation - Remove or replacing an I/O module
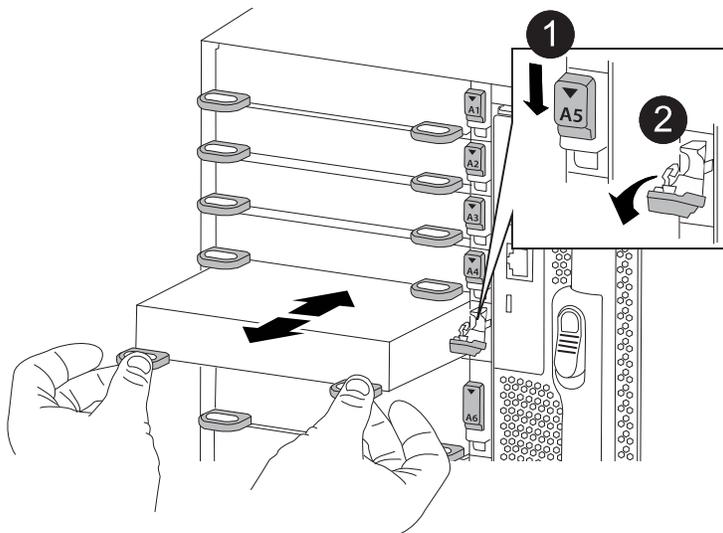
❶	Lettered and numbered I/O cam latch

❷	I/O cam latch completely unlocked

4. Install the I/O module into the target slot:

   a. Align the I/O module with the edges of the slot.

   b. Slide the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.

   c. Push the I/O cam latch all the way up to lock the module in place.

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

6. If the replacement I/O module is a NIC, cable the module or modules to the data switches.

7. Reboot the controller from the LOADER prompt:

   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

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

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

10. If you added:

If I/O module is a…	Then…
NIC module in slots 3 or 7,	Use the `storage port modify -node *<node name> -port *<port name>` `-mode network` command for each port.
Storage module	Install and cable your NS224 shelves, as described in Hot-add workflow.

11. Repeat these steps for controller B.

Replace an LED USB module - ASA A900

The LED USB module provides connectivity to console ports and system status. Replacement of this module does not require tools and does not interrupt service.

Step 1: Replace the LED USB module

Steps

1. Remove the impaired LED USB module:

   Animation - Remove/install LED-USB module
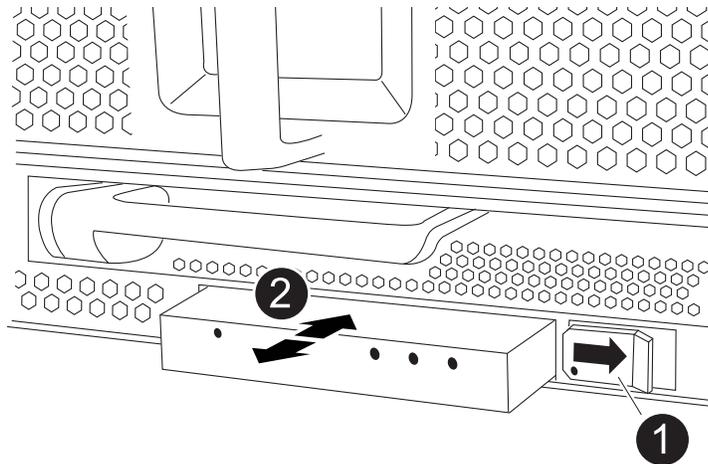
1	Locking button
2	USB LED module

   a. With the bezel removed, locate the LED USB module at the front of the chassis, on the bottom left side.

   b. Slide the latch to partially eject the module.

   c. Pull the module out of the bay to disconnect it from the midplane. Do not leave the slot empty.

2. Install the new LED USB module:

   a. Align the module to the bay with the notch in the corner of the module positioned near the slider latch on the chassis. The bay will prevent you from installing the module upside down.

   b. Push the module into the bay until it is fully seated flush with the chassis.

      There is an audible click when the module is secure and connected to the midplane.

Step 2: Return the failed component

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

Replace the NVRAM module and NVRAM DIMMs - ASA A900

The NVRAM module consists of the NVRAM11 and DIMMs. You can replace a failed NVRAM module or the DIMMs inside the NVRAM module. To replace a failed NVRAM module, you must remove it from the chassis, move the DIMMs to the replacement module, and install the replacement NVRAM module into the chassis.

To replace and NVRAM DIMM, you must remove the NVRAM module from the chassis, replace the failed DIMM in the module, and then reinstall the NVRAM module.

About this task

Because the system ID is derived from the NVRAM module, if replacing the module, disks belonging to the system are reassigned to a new system ID.

Before you begin

- All disk shelves must be working properly.

- If your system is in an HA pair, the partner controller must be able to take over the controller associated with the NVRAM module that is being replaced.

- 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.

- This procedure includes steps for automatically reassigning disks to the controller module associated with the new NVRAM module. You must reassign the disks when directed to in the procedure. Completing the disk reassignment before giveback can cause issues.

- You must replace the failed component with a replacement FRU component you received from your provider.

- You cannot change any disks or disk shelves as part of this procedure.

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

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 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 `y` when prompted.

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

Step 2: Replace the NVRAM module

To replace the NVRAM module, locate it in slot 6 in the chassis and follow the specific sequence of steps.

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

2. Remove the target NVRAM module from the chassis:

   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

   The NVRAM module disengages from the chassis and moves out a few inches.

   c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module face.

   Animation - Replace the NVRAM module

①	Lettered and numbered cam latch

❷	Cam latch completely unlocked

3. Set the NVRAM module on a stable surface and remove the cover from the NVRAM module by pushing down on the blue locking button on the cover, and then, while holding down the blue button, slide the lid off the NVRAM module.

❶	Cover locking button
❷	DIMM and DIMM ejector tabs

4. Remove the DIMMs, one at a time, from the old NVRAM module and install them in the replacement NVRAM module.

5. Close the cover on the module.

6. Install the replacement NVRAM module into the chassis:

   a. Align the module with the edges of the chassis opening in slot 6.

   b. Gently slide the module into the slot until the lettered and numbered cam latch begins to engage with the I/O cam pin, and then push the cam latch all the way up to lock the module in place.

Step 3: Replace a NVRAM DIMM

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

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

2. Remove the target NVRAM module from the chassis:

   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

The NVRAM module disengages from the chassis and moves out a few inches.

c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module face.

[Animation - Replace NVRAM DIMM](#)
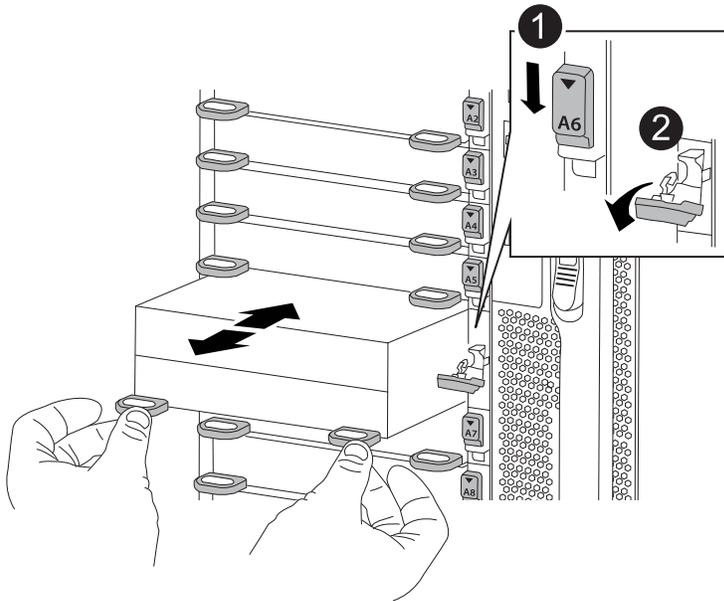
①	Lettered and numbered cam latch
②	cam latch completely unlocked

3. Set the NVRAM module on a stable surface and remove the cover from the NVRAM module by pushing down on the blue locking button on the cover, and then, while holding down the blue button, slide the lid off the NVRAM module.
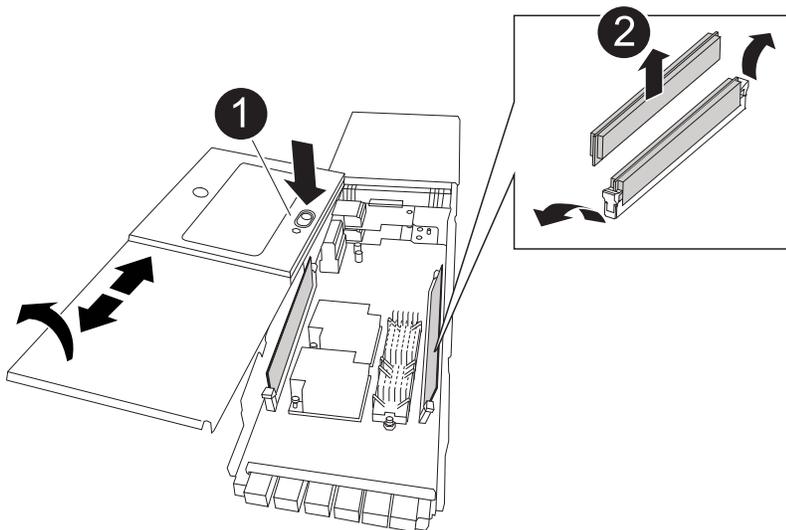
❶	Cover locking button
❷	DIMM and DIMM ejector tabs

4. Locate the DIMM to be replaced inside the NVRAM module, and then remove it by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.

5. 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.

6. Close the cover on the module.

7. Install the NVRAM module into the chassis:

   a. Align the module with the edges of the chassis opening in slot 6.

   b. Gently slide the module into the slot until the lettered and numbered cam latch begins to engage with the I/O cam pin, and then push the cam latch all the way up to lock the module in place.

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 replacement 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 replacement 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 replacement controller, boot the controller and entering y if you are prompted to override the system ID due to a system ID mismatch.

3. Wait until the Waiting for giveback… message is displayed on the console of the controller with the replacement module and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

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

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

4. Give back the controller:

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

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

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

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

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

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

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

```
node1:> storage disk show -ownership

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

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

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

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

   This is required if both of the following are true:

   ◦ The MetroCluster configuration is in a switchover state.

   ◦ The replacement controller is the current owner of the disks on the disaster site.

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

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

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

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

4 entries were displayed.
```

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

10. 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`

Step 6: Return the failed part to NetApp

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

Hot-swap a power supply - ASA A900

Swapping out a power supply involves turning off, disconnecting, and removing the power supply and installing, connecting, and turning on the replacement power supply.

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

About this task
- The power supplies are redundant and hot-swappable. You do not have to shut down the controller to replace a PSU.
- This procedure is written for replacing one power supply at a time.

> (i) It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

- There are four power supplies in the system.
- Power supplies are auto-ranging.

> (i) Do not mix PSUs with different efficiency ratings. Always replace like for like.

Steps
1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.

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

3. Turn off the power supply and disconnect the power cables:

   a. Turn off the power switch on the power supply.

   b. Open the power cable retainer, and then unplug the power cable from the power supply.

4. Press and hold the terra cotta button on the power supply handle, and then pull the power supply out of the chassis.

   CAUTION:

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

   Animation - Remove/install PSU
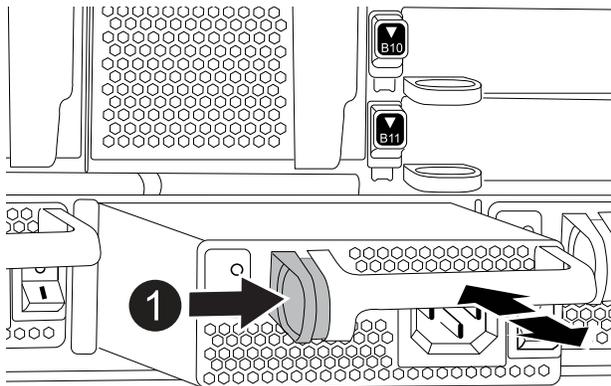
1	Locking button

5. Make sure that the on/off switch of the new power supply is in the Off position.

6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis until it locks into place.

   The power supplies are keyed and can only be installed one way.

   > ⓘ Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Reconnect the power supply cabling:

   a. Reconnect the power cable to the power supply.

   b. Secure the power cable to the power supply using the power cable retainer.

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

8. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

   The green power LED lights when the PSU is fully inserted into the chassis and the amber attention LED flashes initially, but turns off after a few moments.

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

Replace the real-time clock battery - ASA A900

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.

You must use an approved RTC battery.

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

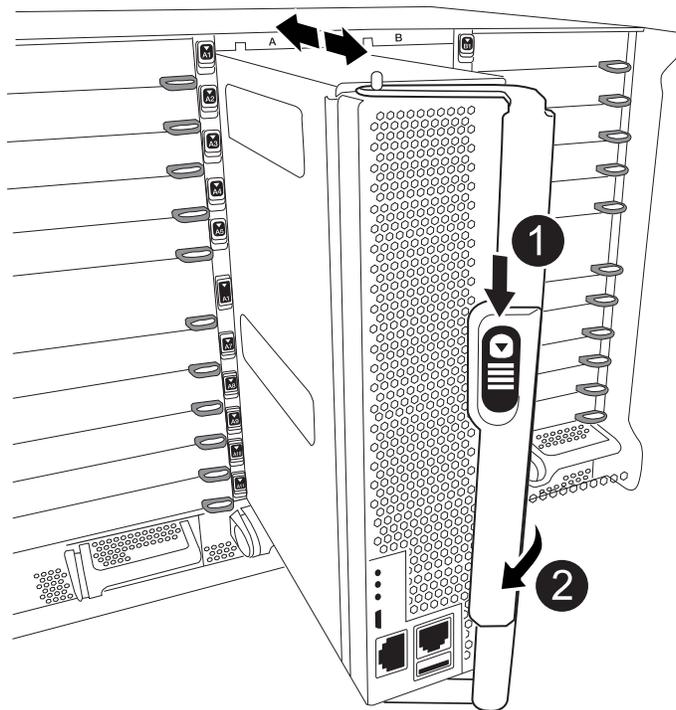| If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

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 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 `y` when prompted.

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

Step 2: Remove the controller

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

Animation - Remove the controller

①	Cam handle release button
②	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

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

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

1	Controller module cover locking button

Step 3: Replace the RTC battery

To replace the RTC battery, you must locate the failed battery in the controller module, remove it from the holder, and then install the replacement battery in the holder.

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

2. Locate the RTC battery.

   Animation - Replace RTC battery
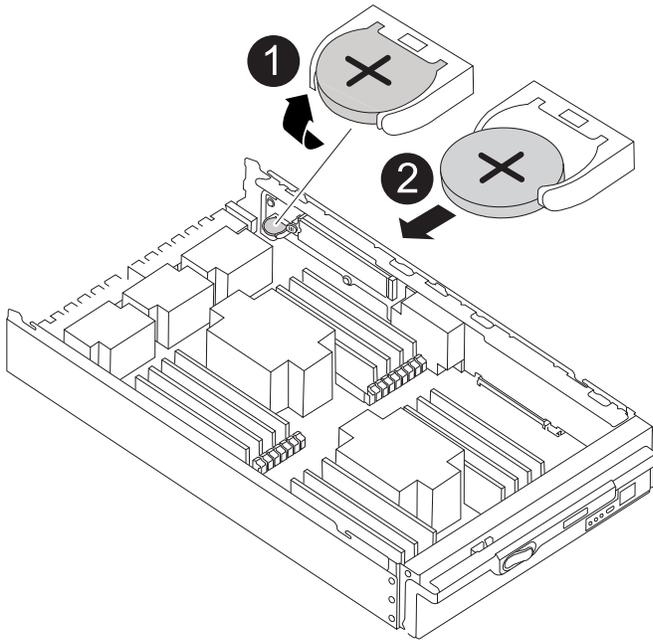
1	RTC battery
2	RTC battery housing

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

> ℹ️ 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. Locate the empty battery holder in the controller module.

6. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.

7. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

8. Reinstall the controller module cover.

Step 4: Reinstall the controller module and set time/date

After you replace the RTC battery, you must reinstall the controller module. If the RTC battery has been left out of the controller module for more than 10 minutes, you may have to reset the time and date.

1. If you have not already done so, close the air duct or controller module cover.

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 media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.

5. Complete the reinstallation of the controller module:

   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

   > (i) Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.

   c. Bind the cables to the cable management device with the hook and loop strap.

   d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.

   e. Halt the controller at the LOADER prompt.

   > (i) 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`.

1. Reset the time and date on the controller:

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

   b. At the LOADER prompt on the target controller, check the time and date.

   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.

   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.

   e. 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 5: Return the failed part to NetApp

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

Key specifications for ASA A900

The following are select specifications for the ASA A900. Visit NetApp Hardware Universe (HWU) for a complete list of ASA A900 specifications. This page is reflective of a single high availability pair.

Key specifications for ASA A900

Platform Configuration: ASA A900 Single Chassis HA Pair

Max Raw Capacity: 14.6880 PB

Memory: 2048.0000 GB

Form Factor: 8U chassis with 2 HA controllers

ONTAP Version: b_startONTAP: 9.16.1P2b_end

PCIe Expansion Slots: 20

Minimum ONTAP Version: ONTAP 9.13.1

Scaleout Maximums

Type	HA Pairs	Raw Capacity	Max Memory
NAS			
SAN	6	88.1 PB / 78.3 PiB	12288 GB
HA Pair		14.7 PB / 13.0 PiB	2048.0000

IO

Onboard IO

No onboard IO data.

Total IO

Protocol	Ports
Ethernet 100 Gbps	32
Ethernet 25 Gbps	64
Ethernet 10 Gbps	64
FC 32 Gbps	64
NVMe/FC 32 Gbps	64
	0
SAS 12 Gbps	64

Management Ports

Protocol	Ports
Ethernet 1 Gbps	2
RS-232 115 Kbps	6
USB 12 Mbps	2

Storage Networking Supported

FC; iSCSI; NVMe/FC ; NVMe/TCP;

System Environment Specifications

- Typical Power: 8004 BTU/hr

- Worst-case Power: 9937 BTU/hr

- Weight: 220.5 lb 100.0 kg

- Height: 8U

- Width: 19" IEC rack-compliant (17.7" 45.0 cm)

- Depth: 28.8" (36.8" with cable management bracket)

- Operating Temp/Altitude/Humidity: 10°C to 35°C (50°F to 95°F) at up to 3048m (10000 ft) elevation;8% to 80% relative humidity, noncondensing

- Non-operating Temp/Humidity: -40°C to 70°C (-40°F to 158°F) up to 12192m (40000 ft) 10% to 95% relative humidity, noncondensing, in original container

- Acoustic Noise: Declared sound power (LwAd): 7.4; Sound pressure (LpAm) (bystander positions): 65.0 dB

Compliance

- Certifications EMC/EMI: AMCA, FCC, ICES, KC, Morocco, VCCI

- Certifications safety: BIS, CB, CSA, G_K_U-SoR, IRAM, NOM, NRCS, SONCAP, TBS

- Certifications Safety/EMC/EMI: EAC, UKRSEPRO

- Certifications Safety/EMC/EMI/RoHS: BSMI, CE DoC, UKCA DoC

- Standards EMC/EMI: BS-EN-55024, BS-EN55035, CISPR 32, EN55022, EN55024, EN55032, EN55035, EN61000-3-2, EN61000-3-3, FCC Part 15 Class A, ICES-003, KS C 9832, KS C 9835

- Standards Safety: ANSI/UL60950-1, ANSI/UL62368-1, BS-EN62368-1, CAN/CSA C22.2 No. 60950-1, CAN/CSA C22.2 No. 62368-1, CNS 14336, EN60825-1, EN62368-1, IEC 62368-1, IEC60950-1, IS 13252(part 1)

High Availability

Ethernet based baseboard management controller (BMC) and ONTAP management interface; Redundant hot-swappable controllers; Redundant hot-swappable power supplies; SAS in-band management over SAS connections;