



AFF A20, AFF A30, and AFF A50 systems

Install and maintain

NetApp
January 09, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap-systems/a20-30-50/overview.html> on January 09, 2026. Always check docs.netapp.com for the latest.

Table of Contents

AFF A20, AFF A30, and AFF A50 systems	1
Key specifications for A20-30-50 platform group	1
AFF A50	1
AFF A20	3
Install and setup	6
Installation and setup workflow - AFF A20, AFF A30, and AFF A50	6
Installation requirements - AFF A20, AFF A30, and AFF A50	7
Prepare to install - AFF A20, AFF A30, and AFF A50	8
Install the hardware - AFF A20, AFF A30, and AFF A50	10
Cable the hardware - AFF A20, AFF A30, and AFF A50	11
Power on the storage system - AFF A20, AFF A30, and AFF A50	20
Maintain	23
Overview of the maintenance procedures - AFF A20, AFF A30, and AFF A50	23
Boot media - automated recovery	24
Boot media - manual recovery	39
Chassis	64
Controller	73
Replace a DIMM - AFF A20, AFF A30, and AFF A50	94
Replace a Drive - AFF A20, AFF A30, and AFF A50	103
Replace a fan module - AFF A20, AFF A30, and AFF A50	106
I/O module	114
Replace the NV battery - AFF A20, AFF A30, and AFF A50	132
Hot-swap a power supply - AFF A20, AFF A30, and AFF A50	140
Replace the real-time clock battery - AFF A20, AFF A30, and AFF A50	144
Key specifications	153
Key specifications for AFF A20	153
Key specifications for AFF A30	155
Key specifications for AFF A50	157

AFF A20, AFF A30, and AFF A50 systems

Key specifications for A20-30-50 platform group

The following are select specifications for the A20-30-50 platform group. Visit [NetApp Hardware Universe](#) (HWU) for a complete list of specifications. This page is reflective of a single high availability pair.

AFF A50

Key specifications for AFF A50

Platform Configuration: AFF A50 Single Chassis HA Pair

Max Raw Capacity: 4.0392 PB

Memory: 256.0000 GB

Form Factor: 2U chassis with 2 HA controllers and 24 drive slots

ONTAP Version: b_startONTAP: 9.16.1P2b_end

PCIe Expansion Slots: 8

Minimum ONTAP Version: ONTAP 9.16.1RC1

Scaleout Maximums

Type	HA Pairs	Raw Capacity	Max Memory
NAS	12	48.5 PB / 43.0 PiB	3072 GB
SAN	6	24.2 PB / 21.5 PiB	1536 GB
HA Pair		4.0 PB / 3.6 PiB	256.0000

IO

Onboard IO

No onboard IO data.

Total IO

Protocol	Ports
Ethernet 100 Gbps	16
Ethernet 25 Gbps	24
Ethernet 10 Gbps	24
FC 64 Gbps	24
NVMe/FC 64 Gbps	24

	0
SAS 12 Gbps	16

Management Ports

Protocol	Ports
Ethernet 1 Gbps	2
RS-232 115 Kbps	4
USB 600 Mbps	2

Storage Networking Supported

CIFS; FC; iSCSI; NFS v3; NFS v4.0; NFS v4.1; NFS v4.2; NVMe/FC ; NVMe/TCP; S3; S3 with NAS; SMB 2.0; SMB 2.1; SMB 2.x; SMB 3.0; SMB 3.1; SMB 3.1.1;

System Environment Specifications

- Typical Power: 3076 BTU/hr
- Worst-case Power: 4860 BTU/hr
- Weight: 60.7 lb 27.5 kg
- Height: 2U
- Width: 19" IEC rack-compliant (17.6" 44.7 cm)
- Depth:
- Operating Temp/Altitude/Humidity: 10°C to 35°C (50°F to 95°F) at up to 3048m (10000 ft) elevation;8% to 80% relative humidity, noncondensing
- Non-operating Temp/Humidity: -40°C to 70°C (-40°F to 158°F) up to 12192m (40000 ft) 10% to 95% relative humidity, noncondensing, in original container
- Acoustic Noise: Declared sound power (LwAd): 8.0; Sound pressure (LpAm) (bystander positions): 70.5 dB

Compliance

- Certifications EMC/EMI: AMCA, FCC, ICES, KC, Morocco, VCCI
- Certifications safety: BIS, CB, CSA, G_K_U-SoR, IRAM, NOM, NRCS, SONCAP, TBS
- Certifications Safety/EMC/EMI: EAC, UKRSEPRO
- Certifications Safety/EMC/EMI/RoHS: BSMI, CE DoC, UKCA DoC
- Standards EMC/EMI: BS-EN-55032, BS-EN55035, CISPR 32, EN55022, EN55024, EN55032, EN55035, EN61000-3-2, EN61000-3-3, FCC Part 15 Class A, ICES-003, KS C 9832, KS C 9835
- Standards Safety: ANSI/UL60950-1, ANSI/UL62368-1, BS-EN62368-1, CAN/CSA C22.2 No. 60950-1, CAN/CSA C22.2 No. 62368-1, CNS 15598-1, EN60825-1, EN62368-1, IEC 62368-1, IEC60950-1, IS 13252(part 1)

High Availability

Ethernet based baseboard management controller (BMC) and ONTAP management interface; Redundant hot-swappable controllers; Redundant hot-swappable power supplies; SAS in-band management over SAS

connections for external shelves;

AFF A20

Key specifications for AFF A20

Platform Configuration: AFF A20 Single Chassis HA Pair

Max Raw Capacity: 4.0392 PB

Memory: 128.0000 GB

Form Factor: 2U chassis with 2 HA controllers and 24 drive slots

ONTAP Version: b_startONTAP: 9.16.1P2b_end

PCIe Expansion Slots: 8

Minimum ONTAP Version: ONTAP 9.16.1RC1

Scaleout Maximums

Type	HA Pairs	Raw Capacity	Max Memory
NAS	3	12.1 PB / 10.8 PiB	384 GB
SAN	3	12.1 PB / 10.8 PiB	384 GB
HA Pair		4.0 PB / 3.6 PiB	128.0000

IO

Onboard IO

No onboard IO data.

Total IO

Protocol	Ports
Ethernet 100 Gbps	12
Ethernet 25 Gbps	32
Ethernet 10 Gbps	24
FC 64 Gbps	24
NVMe/FC 64 Gbps	24
	0
SAS 12 Gbps	16

Management Ports

Protocol	Ports
----------	-------

Ethernet 1 Gbps	2
RS-232 115 Kbps	4
USB 600 Mbps	2

Storage Networking Supported

CIFS; FC; iSCSI; NFS v3; NFS v4.0; NFS v4.1; NFS v4.2; NVMe/FC ; NVMe/TCP; S3; S3 with NAS; SMB 2.0; SMB 2.1; SMB 2.x; SMB 3.0; SMB 3.1; SMB 3.1.1;

System Environment Specifications

- Typical Power: 2040 BTU/hr
- Worst-case Power: 2907 BTU/hr
- Weight: 60.7 lb 27.5 kg
- Height: 2U
- Width: 19" IEC rack-compliant (17.6" 44.7 cm)
- Depth:
- Operating Temp/Altitude/Humidity: 10°C to 35°C (50°F to 95°F) at up to 3048m (10000 ft) elevation;8% to 80% relative humidity, noncondensing
- Non-operating Temp/Humidity: -40°C to 70°C (-40°F to 158°F) up to 12192m (40000 ft) 10% to 95% relative humidity, noncondensing, in original container
- Acoustic Noise: Declared sound power (LwAd): 8.0; Sound pressure (LpAm) (bystander positions): 70.5 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 15598-1, EN60825-1, EN62368-1, IEC 62368-1, IEC60950-1, IS 13252(part 1)

High Availability

Ethernet based baseboard management controller (BMC) and ONTAP management interface; Redundant hot-swappable controllers; Redundant hot-swappable power supplies; SAS in-band management over SAS connections for external shelves;

Key specifications for AFF A30

Platform Configuration: AFF A30 Single Chassis HA Pair

Max Raw Capacity: 4.0392 PB

Memory: 128.0000 GB

Form Factor: 2U chassis with 2 HA controllers and 24 drive slots

ONTAP Version: b_startONTAP: 9.16.1P2b_end

PCIe Expansion Slots: 8

Minimum ONTAP Version: ONTAP 9.16.1RC1

Scaleout Maximums

Type	HA Pairs	Raw Capacity	Max Memory
NAS	4	16.2 PB / 14.3 PiB	512 GB
SAN	4	16.2 PB / 14.3 PiB	512 GB
HA Pair		4.0 PB / 3.6 PiB	128.0000

IO

Onboard IO

No onboard IO data.

Total IO

Protocol	Ports
Ethernet 100 Gbps	16
Ethernet 25 Gbps	24
Ethernet 10 Gbps	24
FC 64 Gbps	24
NVMe/FC 64 Gbps	24
	0
SAS 12 Gbps	16

Management Ports

Protocol	Ports
Ethernet 1 Gbps	2
RS-232 115 Kbps	4
USB 600 Mbps	2

Storage Networking Supported

CIFS; FC; iSCSI; NFS v3; NFS v4.0; NFS v4.1; NFS v4.2; NVMe/FC ; NVMe/TCP; S3; S3 with NAS; SMB 2.0; SMB 2.1; SMB 2.x; SMB 3.0; SMB 3.1; SMB 3.1.1;

System Environment Specifications

- Typical Power: 2892 BTU/hr
- Worst-case Power: 4445 BTU/hr
- Weight: 60.7 lb 27.5 kg
- Height: 2U
- Width: 19" IEC rack-compliant (17.6" 44.7 cm)
- Depth:
- Operating Temp/Altitude/Humidity: 10°C to 35°C 50°F to 95°F) at ≤ 3048m ≤ 10000 ft) elevation; 8% to 80% relative humidity, noncondensing
- Non-operating Temp/Humidity: -40°C to 70°C (-40°F to 158°F) up to 12192m (40000 ft) 10% to 95% relative humidity, noncondensing, in original container
- Acoustic Noise: Declared sound power (LwAd): 8.0; Sound pressure (LpAm) (bystander positions): 70.5 dB

Compliance

- Certifications EMC/EMI: AMCA, FCC, ICES, KC, Morocco, VCCI
- Certifications safety: BIS, CB, CSA, G_K_U-SoR, IRAM, NOM, NRCS, SONCAP, TBS
- Certifications Safety/EMC/EMI: EAC, UKRSEPRO
- Certifications Safety/EMC/EMI/RoHS: BSMI, CE DoC, UKCA DoC
- Standards EMC/EMI: BS-EN-55032, BS-EN55035, CISPR 32, EN55022, EN55024, EN55032, EN55035, EN61000-3-2, EN61000-3-3, FCC Part 15 Class A, ICES-003, KS C 9832, KS C 9835
- Standards Safety: ANSI/UL60950-1, ANSI/UL62368-1, BS-EN62368-1, CAN/CSA C22.2 No. 60950-1, CAN/CSA C22.2 No. 62368-1, CNS 15598-1, EN60825-1, EN62368-1, IEC 62368-1, IEC60950-1, IS 13252(part 1)

High Availability

Ethernet based baseboard management controller (BMC) and ONTAP management interface; Redundant hot-swappable controllers; Redundant hot-swappable power supplies; SAS in-band management over SAS connections for external shelves;

Install and setup

Installation and setup workflow - AFF A20, AFF A30, and AFF A50

To install and set up your AFF A20, AFF A30, or AFF A50 storage system, you must review the installation requirements, prepare your site, install and cable the hardware components, power on the storage system, and set up the ONTAP cluster.



Review the installation requirements

Before installing your storage system, it must meet the installation requirements.

2

Prepare for installation

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

3

Install the hardware

To install the hardware, install the rail kits for your storage system and shelves, and then install and secure your storage system and shelves in the cabinet or telco rack.

4

Cable the hardware

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

5

Power on the storage system

To power on your storage system, power on each shelf and assign a unique shelf ID as needed, and then power on the controllers.

6

Set up your cluster

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

Installation requirements - AFF A20, AFF A30, and AFF A50

Review the requirements for your AFF A20, AFF A30, or AFF A50 storage system.

Equipment needed for install

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

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


Lifting precautions

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

Storage system weight

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


The storage system can weigh up to 61.5 lbs (27.9 kg). To lift the storage system, use two people or a hydraulic lift.



Shelf weight

Take the necessary precautions when moving or lifting your shelf.

An NS224 shelf with NSM100B modules can weigh up to 56.8 lbs (25.8 kg). To lift the shelf, use two people or a hydraulic lift. Keep all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.



Related information

- [Safety information and regulatory notices](#)

What's next?

After you've reviewed the installation requirements and considerations for your storage system, you [prepare for installation](#).

Prepare to install - AFF A20, AFF A30, and AFF A50

Prepare to install your AFF A20, AFF A30, or AFF A50 storage system by getting the site ready, unpacking the boxes and comparing the contents of the boxes to the packing slip, and registering the storage system to access support benefits.

Step 1: Prepare the site

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

Steps

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

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

Step 2: Unpack the boxes

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

Steps

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



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

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

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

Hardware

- Bezel
- Storage system
- Rail kits with instructions (optional)
- Storage shelf (if you ordered additional storage)

Cables

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

Step 3: Register your storage system

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

Steps

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

You can find the serial numbers in the following locations:

- On the packing slip
- In your confirmation email
- On each controller

SSN: XXYYYYYYYYYY



2. Go to the [NetApp Support Site](#).
3. Determine whether you need to register your storage system:

If you are a...	Follow these steps...
Existing NetApp customer	- Sign in with your username and password. - Select Systems > My Systems. - Confirm that the new serial numbers are listed. - If it is not, follow the instructions for new NetApp customers.
New NetApp customer	- Click Register Now, and create an account. - Select Systems > Register Systems. - Enter the storage system's serial numbers and requested details. After your registration is approved, you can download any required software. The approval process might take up to 24 hours.

What's next?

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

Install the hardware - AFF A20, AFF A30, and AFF A50

After you prepare to install your AFF A20, AFF A30, or AFF A50 storage system, install the hardware for the storage system. First, install the rail kits. Then install and secure your storage system in a cabinet or telco rack.

Skip this step if your storage system came in a cabinet.

Before you begin

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

Steps

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

If you are installing multiple shelves, place the first shelf directly above the controllers. Place the

second shelf directly under the controllers. Repeat this pattern for any additional shelves.

- b. Secure the shelf to the cabinet or telco rack using the included mounting screws.

What's next?

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

Cable the hardware - AFF A20, AFF A30, and AFF A50

After you install your AFF A20, AFF A30, or AFF A50 storage system hardware, cable the controllers to the network and shelves.


Before you begin

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

About this task

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

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



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

Step 1: Cable the cluster/HA connections

Create the ONTAP cluster connections. For switchless clusters, connect the controllers to each other. For switched clusters, connect the controllers to the cluster network switches.



The cluster/HA cabling examples show common configurations.

If you do not see your configuration here, go to [NetApp Hardware Universe](#) for comprehensive configuration and slot priority information to cable your storage system.

Switchless cluster cabling

AFF A30 or AFF A50 with two 2-port 40/100 GbE I/O modules

Steps

1. Cable the Cluster/HA interconnect connections:



The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O modules in slots 2 and 4). The ports are 40/100 GbE.

- a. Cable controller A port e2a to controller B port e2a.
- b. Cable controller A port e4a to controller B port e4a.




I/O module ports e2b and e4b are unused and available for host network connectivity.

100 GbE Cluster/HA interconnect cables



Controller A



Controller B

AFF A30 or AFF A50 with one 2-port 40/100 GbE I/O module

Steps

1. Cable the Cluster/HA interconnect connections:

 The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 40/100 GbE.

- a. Cable controller A port e4a to controller B port e4a.
- b. Cable controller A port e4b to controller B port e4b.

100 GbE Cluster/HA interconnect cables



Controller A



Controller B

AFF A20 with one 2-port 10/25 GbE I/O module

Steps

1. Cable the Cluster/HA interconnect connections:

 The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 10/25 GbE.

- a. Cable controller A port e4a to controller B port e4a.
- b. Cable controller A port e4b to controller B port e4b.

25 GbE Cluster/HA interconnect cables



Controller A



Controller B

Switched cluster cabling

AFF A30 or AFF A50 with two 2-port 40/100 GbE I/O modules

Steps

1. Cable the Cluster/HA interconnect connections:




The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O modules in slots 2 and 4). The ports are 40/100 GbE.

- a. Cable controller A port e4a to cluster network switch A.
- b. Cable controller A port e2a to cluster network switch B.
- c. Cable controller B port e4a to cluster network switch A.
- d. Cable controller B port e2a to cluster network switch B.



I/O module ports e2b and e4b are unused and available for host network connectivity.

40/100 GbE Cluster/HA interconnect cables



AFF A30 or AFF A50 with one 2-port 40/100 GbE I/O module

Steps


1. Cable the controllers to the cluster network switches:



The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 40/100 GbE.

- a. Cable controller A port e4a to cluster network switch A.
- b. Cable controller A port e4b to cluster network switch B.
- c. Cable controller B port e4a to cluster network switch A.
- d. Cable controller B port e4b to cluster network switch B.

40/100 GbE Cluster/HA interconnect cables



AFF A20 with one 2-port 10/25 GbE I/O module


1. Cable the controllers to the cluster network switches:



The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 10/25 GbE.

- a. Cable controller A port e4a to cluster network switch A.
- b. Cable controller A port e4b to cluster network switch B.
- c. Cable controller B port e4a to cluster network switch A.
- d. Cable controller B port e4b to cluster network switch B.

10/25 GbE Cluster/HA interconnect cables



Step 2: Cable the host network connections

Cable the controllers to your Ethernet or FC host network.



The host network cabling examples show common configurations.

If you do not see your configuration here, go to [NetApp Hardware Universe](#) for comprehensive configuration and slot priority information to cable your storage system.

Ethernet host cabling

AFF A30 or AFF A50 with two 2-port 40/100 GbE I/O modules


Steps

1. On each controller, cable ports e2b and e4b to the Ethernet host network switches.



The ports on I/O modules in slot 2 and 4 are 40/100 GbE (host connectivity is 40/100 GbE).

40/100 GbE cables




AFF A20, A30 or AFF A50 with one 4-port 10/25 GbE I/O module

Steps

1. On each controller, cable ports e2a, e2b, e2c and e2d to the Ethernet host network switches.

10/25 GbE cables




FC host cabling

AFF A20, A30 or AFF A50 with one 4-port 64 Gb/s FC I/O module

Steps

1. On each controller, cable ports 1a, 1b, 1c and 1d to the FC host network switches.

64 Gb/s FC cables




Step 3: Cable the management network connections

Cable the controllers to your management network.

1. Cable the management (wrench) ports on each controller to the management network switches.

1000BASE-T RJ-45 cables



Do not plug in the power cords yet.

Step 4: Cable the shelf connections

This procedure shows you how to cable the controllers to one NS224 shelf.

About this task

- For the maximum number of shelves supported for your storage system and for all of your cabling options, such as optical and switch-attached, see [NetApp Hardware Universe](#).
- The NS224 shelf cabling procedure shows NSM100B modules instead of NSM100 modules. The cabling is the same regardless of the type of NSM modules used, only the port names are different:
 - NSM100B modules use ports e1a and e1b on an I/O module in slot 1.
 - NSM100 modules use built-in (onboard) ports e0a and e0b.
- You cable each controller to each NSM module on the NS224 shelf using the storage cables that came with your storage system, which could be the following cable type:


100 GbE QSFP28 copper cables



- The graphics show controller A cabling in blue and controller B cabling in yellow.


Steps

1. Cable controller A to the shelf:
 - a. Cable controller A port e3a to NSM A port e1a.
 - b. Cable controller A port e3b to NSM B port e1b.



2. Cable controller B to the shelf:

- a. Cable controller B port e3a to NSM B port e1a.
- b. Cable controller B port e3b to NSM A port e1b.



What's next?

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

Power on the storage system - AFF A20, AFF A30, and AFF A50

After you cable the controllers to the network and shelves in your AFF A20, AFF A30, or AFF A50 storage system, you power on your shelves and controllers.

Step 1: Power on the shelf and assign shelf ID

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

Before you begin

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

About this task

- A valid shelf ID is 01 through 99.

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


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

Steps

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

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

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



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

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

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

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

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

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

The first number continues to blink.

4. Change the second number of the shelf ID:

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

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

The first number on the digital display stops blinking.

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

The second number continues to blink.

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

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

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

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

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

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

7. Replace the left end cap.

Step 2: Power on the controllers

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


Steps

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




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

- b. Using the console cable provided with your storage system, connect one end of the console cable to your laptop and the other end to the serial console port on controller A.
- c. Connect the laptop to the switch on the management subnet.



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



- The system begins to boot. Initial booting might take up to eight minutes.
 - The LEDs flash on and the fans start, which indicates that the controllers are powering on.
 - The fans might be very noisy when they first start up. The fan noise during start-up is normal.
 - The shelf ID display on the front of the system chassis does not illuminate.
4. Secure the power cords using the securing device on each power supply.

What's next?

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

Maintain

Overview of the maintenance procedures - AFF A20, AFF A30, and AFF A50

Maintain the hardware of your AFF A20, AFF A30, or AFF A50 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the AFF A20, AFF A30, or AFF A50 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the AFF A20, AFF A30, or AFF A50 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot the image from a USB drive and restore the configuration from the partner node.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
Drive	A drive is a device that provides the physical storage needed for data.
Fan	A fan cools the controller and drives.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
NV battery	The non-volatile memory (NV) battery is responsible for providing power to the NVMEM components while data in-flight is being destaged to flash memory after a power loss.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - AFF A20, AFF A30, and AFF A50

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 AFF A20, AFF A30, or AFF A50 storage system.

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

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

1

Review the boot media requirements

Review the requirements for boot media replacement.

2

Shut down the controller

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

3

Replace the boot media

Remove the failed boot media from the impaired controller 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 - AFF A20, AFF A30, and AFF A50

Before replacing the boot media in your AFF A20, AFF A30, and AFF A50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0M (wrench) port on the impaired controller is working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

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

Review the following requirements.

- You must replace the failed 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 - AFF A20, AFF A30, and AFF A50

Shut down the impaired controller in your AFF A20, AFF A30, or AFF A50 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

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

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <code>y</code> when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> The <code>-halt true</code> 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 - AFF A20, AFF A30, and AFF A50

The boot media in your AFF A20, AFF A30, or AFF A50 storage system stores essential firmware and configuration data. The replacement process involves removing the controller module, removing the impaired boot media, installing the replacement boot media, and then reinstalling the controller module.

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](#).

About this task

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

A platform chassis has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

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

Step 1: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

Steps


- 1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

- 2. If you are not already grounded, properly ground yourself.
- 3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.


If you are disconnecting a...	Then...
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

- 4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

- 5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Place the controller on an anti-static mat.


7. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Replace the boot media

To replace the boot media, locate it inside the controller and follow the specific sequence of steps.

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

2. Remove the boot media:



1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

3. Install the replacement boot media:


- a. Remove the boot media from its package.
- b. Slide the socket end of the boot media into its socket.
- c. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 3: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



Do not completely insert the controller in the chassis until instructed to do so later in this procedure.

3. Reconnect the cables to the controller; however, do not plug in the power cord to the power supply (PSU) at this time.



Make sure that the console cable is connected to the controller because you want to catch and log the boot sequence later in the boot media replacement procedure when you fully seat the controller in the chassis and it begins to boot.

4. Fully seat the controller in the chassis:

- a. Firmly push on the handles until the controller meets the midplane and is fully seated.

Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.



The controller boots to the LOADER prompt when fully seated in the chassis. It gets its power from the partner controller.

- b. Rotate the controller handles up and lock in place with the tabs.
5. Reconnect the power cord to the PSU on the impaired controller.

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

If you are reconnecting a...	Then...
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

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 - AFF A20, AFF A30, and AFF A50

After installing the new boot media device in your AFF A20, AFF A30, or AFF A50 storage system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

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

Before you begin

- Determine your key manager type:
 - Onboard Key Manager (OKM): Requires cluster-wide passphrase and backup data
 - External Key Manager (EKM): Requires the following files from the partner node:
 - /cfcard/knip/servers.cfg
 - /cfcard/knip/certs/client.crt
 - /cfcard/knip/certs/client.key
 - /cfcard/knip/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: <pre>storage failover giveback -ofnode impaired_node_name</pre> 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.

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

Onboard Key Manager (OKM)

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

```
key manager is configured.  
Entering Bootmenu Option 10...  
  
This option must be used only in disaster recovery procedures. Are  
you sure? (y or n):
```

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

Show example of passphrase and backup data prompts

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

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

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

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

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

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

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

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

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

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

External Key Manager (EKM)

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

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

- a. Enter the EKM configuration settings when prompted:

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

Show example of client certificate contents

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

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

Show example of client key file contents

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

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

Show example of KMIP server file contents

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

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

Show example of server configuration file contents

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

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

Show example of ONTAP Cluster UUID prompt

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

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

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

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

Show example of temporary network setting prompts

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

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

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

b. Verify the key restoration status:

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

Show example of key recovery error and warning messages

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

Uptime: 11m32s
System halting...

LOADER-B>
```

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

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

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

- 5. If automatic giveback was disabled, reenable it:

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

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

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

What's next

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

Return the failed boot media part to NetApp - AFF A20, AFF A30, and AFF A50

If a component in your AFF A20, AFF A30, or AFF A50 storage system fails, return the failed part to NetApp. See the [Part Return and Replacements](#) page for further

information.

Boot media - manual recovery

Boot media manual recovery workflow - AFF A20, AFF A30, and AFF A50

The manual recovery of the boot image involves using a USB drive to reinstall ONTAP onto the AFF A20, AFF A30, or AFF A50 storage system's replacement boot media. You must download the appropriate ONTAP recovery image from the NetApp Support Site and copy it to a USB drive. This prepared USB drive is then used to perform the recovery and restore the system to operational status.

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

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

1

Review the boot media requirements

Review the requirements for replacing the boot media.

2

Check onboard encryption keys

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

3

Shut down the controller

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

4

Replace the boot media

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

5

Boot the recovery image

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

6

Restore encryption

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

7

Return the failed part to NetApp

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

Requirements for manual boot media recovery - AFF A20, AFF A30, and AFF A50

Before replacing the boot media in your AFF A20, AFF A30, or AFF A50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

USB flash drive

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

File preparation

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

Component replacement

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

Controller identification

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

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

What's next?

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

Check encryption support for manual boot media recovery - AFF A20, AFF A30, and AFF A50

To ensure data security on your AFF A20, AFF A30, or AFF A50 storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

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

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

Steps

1. Check if your ONTAP version supports encryption:

```
version -v
```

If the output includes `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



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	<pre>security key-manager keystore show</pre> • 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	<pre>security key-manager show-key-store</pre> • 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.

What's next?

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

Shut down the controller for manual boot media recovery - AFF A20, AFF A30, and AFF A50

Shut down the impaired controller in your AFF A20, AFF A30, or AFF A50 storage system to prevent data loss and maintain system stability during the manual boot media recovery process.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

```
metrocluster node show
```

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

What's next?

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

Replace the boot media and prepare for manual boot recovery - AFF A20, AFF A30, and AFF A50

The boot media in your AFF A20, AFF A30, or AFF A50 storage system stores essential firmware and configuration data. The replacement process involves removing the controller module, removing the impaired boot media, installing the replacement boot media, and then manually transferring the ONTAP image to the replacement boot media using a USB flash drive.

About this task

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

A platform chassis has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

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

Step 1: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

Steps


1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

- 2. If you are not already grounded, properly ground yourself.
- 3. Disconnect the power on the impaired controller:


Power supplies (PSUs) do not have a power switch.

If you are disconnecting a...	Then...
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

- 4. Unplug all cables from the impaired controller.
- Keep track of where the cables were connected.

- 5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Place the controller on an anti-static mat.


7. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Replace the boot media

To replace the boot media, locate it inside the controller and follow the specific sequence of steps.

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

2. Remove the boot media:



1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

3. Install the replacement boot media:


- a. Remove the boot media from its package.
- b. Slide the socket end of the boot media into its socket.
- c. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 3: Reinstall the controller

Reinstall the controller into the chassis, but do not reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



Do not completely insert the controller in the chassis until instructed to do so later in this procedure.

3. Reconnect the cables to the controller; however, do not plug in the power cord to the power supply (PSU) at this time.



Make sure that the console cable is connected to the controller because you want to catch and log the boot sequence later in the boot media replacement procedure when you fully seat the controller in the chassis and it begins to boot.

Step 4: Transfer the boot image to the boot media

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

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- You must have a copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the [Downloads](#) section on the NetApp Support Site

- If NVE is supported, download the image with NetApp Volume Encryption, as indicated in the download button.
- If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.
- You must have a network connection between the node management ports of the controllers (typically the e0M interfaces).

Steps

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



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

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

- c. Remove the USB flash drive from your laptop.
2. Insert the USB flash drive into the USB-A port on the impaired controller.

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

3. Fully seat the impaired controller in the chassis:
 - a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.



The controller boots when fully seated in the chassis. It gets its power from the partner controller.

- b. Rotate the controller handles up and lock in place with the tabs.
4. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

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

5. Reconnect the power cord to the power supply (PSU) on the impaired controller.

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

If you are reconnecting a...	Then...
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.

If you are reconnecting a...	Then...
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next?

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

Manual boot media recovery from a USB drive - AFF A20, AFF A30, and AFF A50

After installing the new boot media device in your AFF A20, AFF A30, or AFF A50 storage system, you can boot the recovery image manually from a USB drive to restore the configuration from the partner node.

Before you begin

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

Steps

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

```
boot_recovery
```

The recovery image is downloaded from the USB flash drive.

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

ONTAP 9.16.0 or earlier

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

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

```
set -privilege advanced
```

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

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



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

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

```
set -privilege admin
```

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

ONTAP 9.16.1 or later

Complete the following steps on the impaired controller:

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

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

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

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

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

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

6. If you disabled automatic giveback, reenable it:

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

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

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

What's next?

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

Restore encryption keys after manual boot recovery - AFF A20, AFF A30, and AFF A50

Restore encryption on the replacement boot media in your AFF A20, AFF A30, or AFF A50 storage system to ensure continued data protection. The replacement process involves verifying key availability, reapplying encryption settings, and confirming secure access to your data.

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

Onboard Key Manager (OKM)

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

Before you begin

Ensure you have the following information available:

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

Steps

On the impaired controller:

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

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. 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. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option
ONTAP 9.7 and earlier	Select the hidden option <code>recover_onboard_keymanager</code> 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. Selection (1-19)? recover_onboard_keymanager </pre>

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

Show example prompt

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

4. Enter the cluster-wide passphrase twice.

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

Show example prompt

```
Enter the passphrase for onboard key management:

Enter the passphrase again to confirm:
```

5. Enter the backup information:
 - a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line, including the dashes.

Show example prompt

Enter the backup data:

-----BEGIN

BACKUP-----

01234567890123456789012345678901234567890123456789012345678901
23

12345678901234567890123456789012345678901234567890123456789012
34

23456789012345678901234567890123456789012345678901234567890123
45

34567890123456789012345678901234567890123456789012345678901234
56

45678901234567890123456789012345678901234567890123456789012345
67

[illegible][illegible][illegible][illegible][illegible][illegible][illegible][illegible][illegible][illegible][illegible][illegible][illegible][illegible][illegible]

```
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
01234567890123456789012345678901234567890123456789012345678901
23
12345678901234567890123456789012345678901234567890123456789012
34
23456789012345678901234567890123456789012345678901234567890123
45
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
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/knip/servers.cfg file or the KMIP server address and port
- /cfcard/knip/certs/client.crt file (client certificate)
- /cfcard/knip/certs/client.key file (client key)
- /cfcard/knip/certs/CA.pem file (KMIP server CA certificates)

Steps

On the impaired controller:

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

Show example boot menu

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

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

Show example prompt

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

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

Show example

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

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

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

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

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

The recovery process completes and displays the following message:

Successfully recovered keymanager secrets.

Show example

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

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

Show example prompt

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

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

6. Restore automatic giveback if you disabled it:

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

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

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

What's next?

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

Return the failed part to NetApp - AFF A20, AFF A30, and AFF A50

If a component in your AFF A20, AFF A30, or AFF A50 storage system fails, return the failed part to NetApp. See the [Part Return and Replacements](#) page for further information.

Chassis

Chassis replacement workflow - AFF A20, AFF A30, and AFF A50

Get started with replacing the chassis of your AFF A20, AFF A30, or AFF A50 storage

system by reviewing the replacement requirements, shutting down the controllers, replacing the chassis, and verifying system operations.

1

Review the chassis replace requirements

To replace the chassis, you must meet certain requirements.

2

Shut down the controllers

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

3

Replace the chassis

Replacing the chassis includes moving the drives and any drive blanks, controllers (with the power supplies), and bezel from the impaired chassis to the new chassis, and swapping out the impaired chassis with the new chassis of the same model as the impaired chassis.

4

Complete chassis replacement

Verify the HA state of the chassis and return the failed part to NetApp.

Requirements to replace the chassis - AFF A20, AFF A30, and AFF A50

Before replacing the chassis of your AFF A20, AFF A30, or AFF A50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement chassis, and the necessary tools.

Review the following requirements and considerations.

Requirements

- The replacement chassis must be the same model as the impaired chassis. This procedure is for a like-for-like replacement, not for an upgrade.
- All other components in the storage system must be functioning properly; if not, contact [NetApp Support](#) before continuing with this procedure.

Considerations

- The chassis replacement procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.
- You can use the chassis replacement procedure with all versions of ONTAP supported by your storage system.
- The chassis replacement procedure is written with the assumption that you are moving the bezel, drives, any drive blanks, and controllers to the new chassis.

What's next?

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

Shut down the controllers to replace the chassis - AFF A20, AFF A30, and AFF A50

Shut down the controllers in your AFF A20, AFF A30, or AFF A50 storage system to prevent data loss and ensure system stability when replacing the chassis.

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

Before you begin

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

Steps

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

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

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

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

6. Exit the cluster shell:

```
exit
```

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

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

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

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



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

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

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

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

What's next?

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

Replace the chassis - AFF A20, AFF A30, and AFF A50

Replace the chassis of your AFF A20, AFF A30, or AFF A50 storage system when a hardware failure requires it. The replacement process involves removing the controllers, removing the drives, installing the replacement chassis, and reinstalling the chassis components.

About this task

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

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

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

Step 1: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

Steps


1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

- If you are not already grounded, properly ground yourself.
- Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.


If you are disconnecting a...	Then...
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

- Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

- Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Repeat these steps for the other controller in the chassis.


Step 2: Remove the drives from the impaired chassis

You need to remove all of the drives and any drive blanks from the impaired chassis so that later in the procedure you can install them in the replacement chassis.

1. Gently remove the bezel from the front of the storage system.
2. Remove the drives and any drive blanks:



Keep track of what drive bay each drive and drive blank was removed from because they must be installed in the same drive bays in the replacement chassis.



1	Press the release button on the drive face to open the cam handle.
2	Rotate the cam handle downward to disengage the drive from the midplane.

3	Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand. When removing a drive, always use two hands to support its weight.  Because drives are fragile, minimize handling to avoid damaging them.
-----------------	--

3. Set the drives aside on a static-free cart or table.

Step 2: Replace the chassis from within the equipment rack or system cabinet

You remove the impaired chassis from the equipment rack or system cabinet, install the replacement chassis, install the drives, any drive blanks, and then install the bezel.

1. Remove the screws from the impaired chassis mount points.

Set the screws aside to use later in this procedure.

 If the storage system shipped in a NetApp system cabinet, you must remove additional screws at the rear of the chassis before the chassis can be removed.


- Using two people or a power lift, remove the impaired chassis from the equipment rack or system cabinet by sliding it off the rails, and then set it aside.
- Using two people, install the replacement chassis into the equipment rack or system cabinet by sliding it onto the rails.
- Secure the front of the replacement chassis to the equipment rack or system cabinet using the screws you removed from the impaired chassis.

Step 4: Install the controllers and drives

Install the controllers and drives into the replacement chassis and reboot the controllers.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when installing a controller, and can be used as a reference for the rest of the controller installation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis and push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

1. Insert one of the controllers into the chassis:

- a. Align the back of the controller with the opening in the chassis.
- b. Firmly push on the handles until the controller meets the midplane and is fully seated in the chassis.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

- c. Rotate the controller handles up and lock in place with the tabs.

2. Recable the controller, as needed, except for the power cords.

3. Repeat these steps to install the second controller into the chassis.

4. Install the drives and any drive blanks you removed from the impaired chassis into the replacement chassis:



The drives and drive blanks must be installed in the same drive bays in the replacement chassis.

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

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

- d. Repeat the process for the remaining drives.

5. Install the bezel.

6. Reconnect the power cords to the power supplies (PSU) in the controllers.

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



The controllers begin to boot as soon as the power is restored.

If you are reconnecting a...	Then...
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.

If you are reconnecting a...	Then...
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. If controllers boot to the LOADER prompt, reboot the controllers:

```
boot_ontap
```

8. Turn AutoSupport back on:

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

What's next?

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

Complete chassis replacement - AFF A20, AFF A30, and AFF A50

Verify the HA state of the chassis and then return the failed part to NetApp to complete the final step in the AFF A20, AFF A30, and AFF A50 chassis replacement procedure.

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

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

1. In Maintenance mode, from either controller, display the HA state of the local controller and chassis:

```
ha-config show
```

The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your storage system configuration:
 - a. Set the HA state for the chassis:

```
ha-config modify chassis HA-state
```

The value for HA-state should be *ha*. The value for HA-state can be one of the following: * **ha** * *mcc* (not supported in ASA)

- b. Confirm that the setting has changed:

```
ha-config show
```

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

Step 2: Return the failed part to NetApp

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

Controller

Controller replacement workflow - AFF A20, AFF A30, and AFF A50

Get started with replacing the controller in your AFF A20, AFF A30, or AFF A50 storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.

1

Review the controller replacement requirements

To replace the controller, you must meet certain requirements.

2

Shut down the impaired controller

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

3

Replace the controller

Replacing the controller includes removing the impaired controller, moving FRU components to the replacement controller, installing the replacement controller in the chassis, setting the time and date, and then recabling.

4

Restore and verify the system configuration

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

5

Give back the controller

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

6

Complete controller replacement

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

Requirements to replace the controller - AFF A20, AFF A30, and AFF A50

Before replacing the controller in your AFF A20, AFF A30, or AFF A50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

Review the requirements and considerations for the controller replacement procedure.

Requirements

- All shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section [Choosing the correct recovery procedure](#) to determine whether you should use this procedure.
- You must replace a controller with a controller of the same model type. You cannot upgrade your system by just replacing the controller.
- You cannot change any drives or shelves as part of this procedure.
- You must always capture the controller’s console output to a text log file.

The console output provides you with a record of the procedure you can use to troubleshoot issues you might encounter during the replacement process.

Considerations

It is important that you apply the commands in this procedure to the correct controller:

- The *impaired* controller is the controller that is being replaced.
- The *replacement* controller is the new controller that is replacing the impaired controller.
- The *healthy* controller is the surviving controller.

What’s next?

After you’ve reviewed the requirements to replace the impaired controller, you need to [shut down the impaired controller](#).

Shut down the impaired controller - AFF A20, AFF A30, and AFF A50

Shut down the impaired controller in your AFF A20, AFF A30, or AFF A50 storage system to prevent data loss and ensure system stability when replacing the controller.

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

```
metrocluster node show
```

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

What's next?

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

Replace the controller - AFF A20, AFF A30, and AFF A50

Replace the controller in your AFF A20, AFF A30, or AFF A50 storage system when a hardware failure requires it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

About this task

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

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

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

Step 1: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

Steps


1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

- 2. If you are not already grounded, properly ground yourself.
- 3. Disconnect the power on the impaired controller:


Power supplies (PSUs) do not have a power switch.

If you are disconnecting a...	Then...
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

- 4. Unplug all cables from the impaired controller.
- Keep track of where the cables were connected.

- 5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Move the power supply

Move the power supply (PSU) to the replacement controller.

1. Move the PSU from the impaired controller:


Make sure the left side controller handle is in the upright position to allow you access to the PSU.

Option 1: Move an AC PSU

To move an AC PSU, complete the following steps.

Steps

1. Remove the AC PSU from the impaired controller:



1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	 The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you. Pull the PSU out of the controller while using your other hand to support its weight.

2. Insert the PSU into the replacement controller:

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

A PSU will only properly engage with the internal connector and lock in place one way.



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

- c. Rotate the handle down, so it is out of the way of normal operations.

Option 2: Move a DC PSU

To move a DC PSU, complete the following steps.


Steps

1. Remove the DC PSU from the impaired controller:

- a. Rotate the handle up, to its horizontal position, and then grasp it.
- b. With your thumb, press the terracotta tab to release the locking mechanism.
- c. Pull the PSU out of the controller while using your other hand to support its weight.



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



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

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

A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.




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

- c. Rotate the handle down, so it is out of the way of normal operations.

Step 3: Move the fans

Move the fans to the replacement controller.

1. Remove one of the fans from the impaired controller:




1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

2. Insert the fan into the replacement controller by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.
3. Repeat these steps for the remaining fans.

Step 4: Move the NV battery

Move the NV battery to the replacement controller.

1. Remove the NV battery from the impaired controller:



1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	1. Push in and hold the tab on the connector. 2. Pull the connector up and out of the socket. As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.

2. Install the NV battery into the replacement controller:

- Plug the wiring connector into its socket.
- Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
- Place the NV battery into the compartment.


The NV battery should sit flush in its compartment.

Step 5: Move system DIMMs

Move the DIMMs to the replacement controller.

If you have DIMM blanks, you do not need to move them, the replacement controller should come with them installed.

1. Remove one of the DIMMs from the impaired controller:



1	DIMM slot numbering and positions.  Depending on your storage system model, you will have two or four DIMMs.
2	• Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller in the proper orientation. • Eject the DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot.  Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

2. Install the DIMM in the replacement controller:

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

The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

When inserted correctly, the DIMM goes in easily but fits tightly in the slot. If not, reinsert the DIMM.


- Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
- Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.

3. Repeat these steps for the remaining DIMMs.

Step 6: Move the boot media

Move the boot media to the replacement controller.

1. Remove the boot media from the impaired controller:



1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

2. Install the boot media into the replacement controller:

- a. Slide the socket end of the boot media into its socket.
- b. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 7: Move the I/O modules

Move the I/O modules and any I/O blanking modules to the replacement controller.


1. Unplug cabling from one of the I/O modules.

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

2. Remove the I/O module from the impaired controller:

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

If you are removing the I/O module in slot 4, make sure the right side controller handle is in the upright position to allow you access to the I/O module.



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

3. Install the I/O module into the replacement controller:

- a. Align the I/O module with the edges of the slot.
- b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O Module.

- c. Turn the thumbscrew clockwise to tighten.


4. Repeat these steps to move the remaining I/O modules and any I/O blanking modules to the replacement controller.

Step 8: Install the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



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

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.



Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
 - a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

- b. Rotate the controller handles up and lock in place with the tabs.



The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Take the controller to the LOADER prompt by pressing CTRL-C to abort AUTOBOOT.
6. Set the time and date on the controller:

Make sure you are at the controller's LOADER prompt.

- a. Display the date and time on the controller:

```
show date
```



Time and date default is in GMT. You have the option to display in local time and in 24hr mode.

- b. Set the current time in GMT:

```
set time hh:mm:ss
```

You can get the current GMT from the healthy node:

```
date -u
```

- c. Set the current date in GMT:

```
set date mm/dd/yyyy
```

You can get the current GMT from the healthy node:

```
date -u
```

7. Recable the controller as needed.
8. Reconnect the power cord to the power supply (PSU).

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

If you are reconnecting a...	Then...
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next?

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

Restore and verify the system configuration - AFF A20, AFF A30, and AFF A50

Verify that the controller's HA configuration is active and functioning correctly in your AFF A20, AFF A30, or AFF A50 storage system, and confirm that the system's adapters list all the paths to the disks.

Step 1: Verify HA config settings

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

1. Boot to maintenance mode:

```
boot_ontap maint
```

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

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

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



If you see *PERSONALITY MISMATCH* contact customer support.

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

```
ha-config show
```

The HA state should be the same for all components.

5. If the displayed system state of the controller does not match your storage system configuration, set the HA state for the controller:

```
ha-config modify controller ha
```

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

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

6. Confirm that the setting has changed:

```
ha-config show
```

Step 2: Verify disk list

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

```
storage show disk -p
```

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

2. Exit Maintenance mode:

halt

What's next?

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

Give back the controller - AFF A20, AFF A30, and AFF A50

Return control of storage resources to the replacement controller so your AFF A20, AFF A30, or AFF A50 storage system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption, Onboard Key Manager (OKM) encryption, or External Key Manager (EKM) encryption.

No encryption

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

Steps

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

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

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



You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option 1 for normal boot.
6. Press <enter> when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as `admin`.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`



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

9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
 - a. Move the console cable back to the replacement controller.
 - b. Synchronize missing keys: `security key-manager onboard sync`



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

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

The output should show no results when properly synchronized.

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

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

External key manager (EKM)

Reset encryption and return the controller to normal operation.

Steps

1. If the root volume is encrypted with External Key Manager and the console cable is connected to the replacement node, enter `boot_ontap` menu and select option 11.
2. If these questions appear, answer `y` or `n` as appropriate:

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

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

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

Do you have a copy of the `/cfcard/kmip/servers.cfg` file? {y/n}

Do you know the KMIP server address? {y/n}

Do you know the KMIP port? {y/n}



Contact [NetApp Support](#) if you have issues.

3. Supply the information for:
 - The client certificate (`client.crt`) file contents
 - The client key (`client.key`) file contents
 - The KMIP server CA(s) (`CA.pem`) file contents
 - The IP address for the KMIP server
 - The port for the KMIP server
4. Once the system processes, you see the Boot Menu. Select '1' for normal boot.
5. Check the takeover status: `storage failover show`
6. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
7. If automatic giveback was disabled, reenable it: `storage failover modify -node local`

```
-auto-giveback true
```

8. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

What's next?

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

Complete controller replacement - AFF A20, AFF A30, and AFF A50

To complete the controller replacement for your AFF A20, AFF A30, or AFF A50 storage system, first restore the NetApp Storage Encryption configuration (if necessary) and install the required licenses on the new controller. Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, register the new controller's serial number and then return the failed part to NetApp.

Step 1: Install licenses for the replacement controller in ONTAP

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

Before you begin

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

About this task

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

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

- The 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, register the serial number, and check cluster health

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

Steps

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

If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`
2. Register the system serial number with NetApp Support.
 - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
 - If AutoSupport is not enabled, call [NetApp Support](#) to register the serial number.
3. Check the health of your cluster. See the [How to perform a cluster health check with a script in ONTAP KB](#) article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 3: Return the failed part to NetApp

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

Replace a DIMM - AFF A20, AFF A30, and AFF A50

Replace a DIMM in your AFF A20, AFF A30, or AFF A50 storage system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

Before you begin

- All other components in the storage system must be working correctly; if not, contact [NetApp Support](#) before continuing.
- You must replace the failed FRU component with a replacement FRU component you received from your provider.

About this task

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

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

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

Step 1: Shut down the impaired controller

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 <code>y</code> 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: <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> The <i>-halt true</i> parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

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

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

```
metrocluster node show
```

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Step 2: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

Steps


1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.


If you are disconnecting a...	Then...
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace a DIMM

To replace a DIMM, locate the faulty DIMM inside the controller and follow the specific sequence of steps.


Steps

- 1. If you are not already grounded, properly ground yourself.
- 2. Locate the DIMMs on your controller and identify the faulty DIMM.



Consult either the [Netapp Hardware Universe](#) or the FRU map on the cover of the controller for exact DIMM locations.

- 3. Remove the faulty DIMM:



1	DIMM slot numbering and positions.  Depending on your storage system model you will have two or four DIMMs.
2	• Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM using the same orientation. • Eject the faulty DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot.  Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

4. Install the replacement DIMM:

- a. Remove the replacement DIMM from its antistatic shipping bag.
- b. Make sure that the DIMM ejector tabs on the connector are in the open position.
- c. Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

When inserted correctly, the DIMM goes in easily but fits tightly in the slot. Reinsert the DIMM if you feel it is not inserted correctly.


- d. Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
- e. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



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

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.



Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
 - a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

- b. Rotate the controller handles up and lock in place with the tabs.



The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

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

If you are reconnecting a...	Then...
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

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

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

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

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

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

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


Step 5: Return the failed part to NetApp

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

Replace a Drive - AFF A20, AFF A30, and AFF A50

Replace a drive in your AFF A20, AFF A30, or AFF A50 storage system when a drive fails or requires an upgrade. The replacement process involves identifying the faulty drive, safely removing it, and installing a new drive to ensure continued data access and system performance.

You can replace a failed SSD drive nondisruptively while I/O is in progress.

Before you begin

- The drive that you are installing must be supported by your storage system.

[NetApp Hardware Universe](#)

- If self-encrypting drive (SED) authentication is enabled, you must use the SED replacement instructions in the ONTAP documentation.

Instructions in the ONTAP documentation describe additional steps you must perform before and after replacing an SED.

[NetApp encryption overview with the CLI](#)

- All other components in the storage system must be functioning properly; if not, contact [NetApp Support](#) before continuing with this procedure.
- Verify that the drive you are removing is failed.

You can verify that the drive is failed by running the `storage disk show -broken` command. The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.



Depending on the drive type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

About this task

- When replacing a failed drive, you must wait 70 seconds between the removal of the drive and the insertion of the replacement drive to allow the storage system to recognize that a drive was removed.
- The best practice is to have the current version of the Disk Qualification Package (DQP) installed before hot-swapping a drive.

Having the current version of the DQP installed allows your system to recognize and use newly qualified drives. This avoids system event messages about having noncurrent drive information and prevention of drive partitioning because drives are not recognized. The DQP also notifies you of noncurrent drive firmware.

[NetApp Downloads: Disk Qualification Package](#)

- The best practice is to have current versions of NVMe shelf module (NSM) firmware and drive firmware on your system before replacing FRU components.

[NetApp Downloads: Disk Shelf Firmware](#)

[NetApp Downloads: Disk Drive Firmware](#)



Do not revert firmware to a version that does not support your shelf and its components.

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



Drive firmware checks occur every two minutes.

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

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

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

Steps

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



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

- a. Verify whether automatic drive assignment is enabled:

```
storage disk option show
```

You can enter the command on either controller.

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

- b. If automatic drive assignment is enabled, disable it:

```
storage disk option modify -node node_name -autoassign off
```

You must disable automatic drive assignment on both controllers.

2. Properly ground yourself.
3. Remove the bezel from the front of the storage system.
4. Physically identify the failed drive.


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

illuminate.



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

5. Remove the failed drive:



1	Press the release button on the drive face to open the cam handle.
2	Rotate the cam handle downward to disengage the drive from the midplane.
3	Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand. When removing a drive, always use two hands to support its weight. Because drives are fragile, minimize handling to avoid damaging them.

6. Wait a minimum of 70 seconds before inserting the replacement drive.

7. Insert the replacement drive:

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

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

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

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

updating, the LED blinks.

9. If you are replacing another drive, repeat the preceding steps.
10. Reinstall the bezel on the front of the storage system.
11. If you disabled automatic drive assignment earlier in this procedure, manually assign drive ownership, and then reenable automatic drive assignment if needed:
 - a. Display all unowned drives:

```
storage disk show -container-type unassigned
```

You can enter the command on either controller.

- b. Assign each drive:

```
storage disk assign -disk disk_name -owner owner_name
```

You can enter the command on either controller.

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

- c. Reenable automatic drive assignment if needed:

```
storage disk option modify -node node_name -autoassign on
```

You must reenable automatic drive assignment on both controllers.

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

Contact technical support at [NetApp Support](#), 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure. [//2025-11-17 ontap-systems-internal/issues/1391](#)

Replace a fan module - AFF A20, AFF A30, and AFF A50

Replace a fan module in your AFF A20, AFF A30, or AFF A50 storage system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

About this task

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

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

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

Step 1: Shut down the impaired controller

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 <code>y</code> 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: <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> The <i>-halt true</i> parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

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

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

```
metrocluster node show
```

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Step 2: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

Steps


1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.


If you are disconnecting a...	Then...
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.


6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace fan

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

Steps

- 1. Identify the fan that you must replace by checking the console error messages.
- 2. Remove the failed fan:



1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.


- 3. Insert the replacement fan by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.

Step 4: Reinstall the controller module

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



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

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.



Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
 - a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

- b. Rotate the controller handles up and lock in place with the tabs.



The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

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

If you are reconnecting a...	Then...
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

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

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

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

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

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

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

Step 5: Return the failed part to NetApp

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

I/O module

Overview of I/O module maintenance - AFF A20, AFF A30, and AFF A50

The AFF A20, AFF A30, and AFF A50 storage systems offer flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding, hot-swapping, or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your storage system with the same type of I/O module, or with a different type of I/O module. You can hot-swap a cluster and HA I/O module when your storage system meets specific requirements. You can also add an I/O module to a storage system with available slots.

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

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

- [Hot-swap a cluster and HA I/O module](#)

Hot-swapping a failed cluster and HA I/O module can restore the storage system to its optimal operating state. Hot-swapping is done without having to manually take over the impaired controller.

To use this procedure, your storage system must be running ONTAP 9.17.1 or later and meet specific storage system requirements.

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

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

Add an I/O module - AFF A20, AFF A30, and AFF A50

Add an I/O module to your AFF A20, AFF A30, or AFF A50 storage system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your AFF A20, AFF A30, and AFF A50 storage systems when there are slots available or when all slots are fully populated.

About this task

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

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

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

Step 1: Shut down the impaired controller module

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: 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](#).
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

```
metrocluster node show
```

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Step 2: Add the new I/O module

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

Before you begin

- Check the [NetApp Hardware Universe](#) to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in [NetApp Hardware Universe](#) and use the best one available for your I/O module.
- All other components in the storage system must be functioning properly; if not, contact [NetApp Support](#) before continuing with this procedure.


Add I/O module to an available slot

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

Steps

1. If you are not already grounded, properly ground yourself.
2. On the impaired controller, remove the I/O blanking module from the target slot.

Unused I/O slots should have blanking module installed to prevent possible thermal issues and for EMC compliance.



1	On the I/O blanking module, turn the thumbscrew counterclockwise to loosen.
2	Pull the I/O blanking module out of the controller using the tab on the left and the thumbscrew.

3. Install the new I/O module:
 - a. Align the I/O module with the edges of the controller slot opening.
 - b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O module.

- c. Turn the thumbscrew clockwise to tighten.
4. Cable the I/O module to the designated devices.

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

5. Reboot the impaired controller from the LOADER prompt: `bye`

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

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

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

7. Repeat these steps to add an I/O module to the other controller.

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

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

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

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

Add I/O module to a fully-populated system

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


About this task

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

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See 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. On the impaired controller, unplug any cabling on the target I/O module.
3. Remove the target I/O module from the controller:



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

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

- a. Align the I/O module with the edges of the slot.
- b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O Module.

- c. Turn the thumbscrew clockwise to tighten.

5. Cable the I/O module to the designated devices.

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

6. Repeat the I/O module remove and install steps to add any additional I/O modules in the controller.

7. Reboot the impaired controller from the LOADER prompt:

```
bye
```

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

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

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

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

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

10. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

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

11. If you installed a NIC module, specify the usage mode for each port as *network*:

```
storage port modify -node node_name -port port_name -mode network
```

12. Repeat these steps for the other controller.

Hot-swap an I/O module used for cluster and HA traffic - AFF A20, AFF A30, and AFF A50

The cluster and HA I/O module supports interconnects for clustering and high-availability. You can hot-swap the module in your AFF A20, AFF A30, or AFF A50 storage system when the module fails and if your storage system meets specific requirements.

To hot-swap a module, you ensure your storage system meets the procedure requirements, prepare the storage system and I/O module in slot 4, hot-swap the failed module for an equivalent one, bring the replacement module online, restore the storage system to normal operation, and return the failed module to NetApp.

About this task

- Hot-swapping the cluster and HA I/O module means that you do not have to perform a manual takeover; the impaired controller (the controller with the failed cluster and HA I/O module) has automatically taken over the healthy controller.

When the impaired controller has taken over the healthy controller, the only way to recover without an outage is to hot-swap the module.

- It is critical to apply the commands to the correct controller when you are hot-swapping the cluster and HA I/O module:
 - The *impaired controller* is the controller on which you are hot-swapping the cluster and HA I/O module and it is the controller that has taken over the healthy controller.
 - The *healthy controller* is the HA partner of the impaired controller and it is the controller that was taken over by the impaired controller.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

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

Step 1: Ensure the storage system meets the procedure requirements

To use this procedure, make sure your storage system meets all requirements.



If your storage system does not meet all requirements, you must use the [replace an I/O module procedure](#).

- Your storage system must be running ONTAP 9.17.1 or later.
- The I/O module that failed must be a cluster and HA I/O module in slot 4 and you must be replacing it with an equivalent cluster and HA I/O module. You cannot change the I/O module type.
- Your storage system configuration must have only one cluster and HA I/O module located in slot 4, not two cluster and HA I/O modules.
- Your storage system must be a two-node (switchless or switched) cluster configuration.
- The controller with the failed cluster and HA I/O module (the impaired controller) must have already taken over the healthy partner controller. The takeover should have occurred automatically if the I/O module is failed.

For two-node clusters, the storage system cannot discern which controller has the failed I/O module, so either controller might initiate the takeover. The cluster and HA I/O module hot-swap procedure is only supported when the controller with the failed I/O module (the impaired controller) has taken over the healthy controller.

You can verify that the impaired controller successfully took over the healthy controller by entering the `storage failover show` command.

If you are not sure which controller has the failed I/O module, contact [NetApp Support](#).

- All other components in the storage system must be functioning properly; if not, contact [NetApp Support](#) before continuing with this procedure.

Step 2: Prepare the storage system and I/O module slot 4

Prepare the storage system and I/O module slot 4 so that it is safe to remove the failed cluster and HA I/O module:

Steps

1. Properly ground yourself.
2. Unplug cabling from the failed cluster and HA I/O module.

Make sure to label the cables so that later in this procedure you can reconnect them to the same ports.

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

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

For example, the following AutoSupport message suppresses automatic case creation for two hours:

```
node2::> system node autosupport invoke -node * -type all -message MAINT=2h
```

4. Disable automatic giveback:
 - a. Enter the following command from the console of the impaired controller:
5. Prepare the failed cluster and HA module in slot 4 for removal by removing it from service and powering it off:
 - a. Enter the following command:

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

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

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

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

For example, the following command prepares the module in slot 4 on node 2 (the impaired controller) for removal, and displays a message that it is safe to remove:

```
node2::> system controller slot module remove -node node2 -slot 4
```

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be powered off for removal.

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

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

6. Verify the failed cluster and HA module in slot 4 is powered off:

```
system controller slot module show
```


The output should show *powered-off* in the status column for the failed module in slot 4.

Step 3: Replace the failed cluster and HA I/O module

Replace the failed cluster and HA I/O module in slot 4 with an equivalent I/O module:

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the failed cluster and HA I/O module from the impaired controller:



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew on the right.

3. Install the replacement cluster and HA I/O module into slot 4:

- a. Align the I/O module with the edges of the slot.
- b. Gently push the I/O module all the way into the slot, making sure to properly seat the I/O module into the connector.

You can use the tab on the left and the thumbscrew on the right to push in the I/O module.

- c. Turn the thumbscrew clockwise to tighten.

4. Cable the cluster and HA I/O module.

Step 4: Bring the replacement cluster and HA I/O module online

Bring the replacement cluster and HA I/O module in slot 4 online, verify the module ports initialized successfully, verify slot 4 is powered on, and then verify the module is online and recognized.

Steps

1. Bring the replacement cluster and HA I/O module online:

- a. Enter the following command:

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

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

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

For example, the following command brings slot 4 on node 2 (the impaired controller) online, and displays a message that the process was successful:

```
node2::> system controller slot module insert -node node2 -slot 4  
  
Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be  
powered on and initialized.  
  
Do you want to continue? {y|n}: `y`  
  
The module has been successfully powered on, initialized and placed  
into service.
```

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

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



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

The output should show a `hotplug.init.success` EMS event logged for each port on the cluster and HA I/O module with `hotplug.init.success`: in the *Event* column.

For example, the following output shows initialization succeeded for cluster and HA I/O module ports e4b and e4a:

```
node2::> event log show -event *hotplug.init*

Time                Node                Severity      Event
-----
-----

7/11/2025 16:04:06  node2          NOTICE      hotplug.init.success:
Initialization of ports "e4b" in slot 4 succeeded

7/11/2025 16:04:06  node2          NOTICE      hotplug.init.success:
Initialization of ports "e4a" in slot 4 succeeded

2 entries were displayed.
```

3. Verify I/O module slot 4 is powered on and ready for operation:

```
system controller slot module show
```

The output should show slot 4 status as *powered-on* and therefore ready for operation of the replacement cluster and HA I/O module.

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

Enter the command from the console of the impaired controller:

```
system controller config show -node local -slot4
```

If the replacement cluster and HA I/O module was successfully brought online and is recognize, the output shows I/O module information, including port information, for slot 4.

For example, you should see output similar to the following:


```

node2::> system controller config show -node local -slot 4

Node: node2
Sub- Device/
Slot slot Information
-----
  4      - Dual 40G/100G Ethernet Controller CX6-DX
           e4a MAC Address: d0:39:ea:59:69:74 (auto-100g_cr4-fd-
up)
           QSFP Vendor:          CISCO-BIZLINK
           QSFP Part Number:     L45593-D218-D10
           QSFP Serial Number:   LCC2807GJFM-B
           e4b MAC Address: d0:39:ea:59:69:75 (auto-100g_cr4-fd-
up)
           QSFP Vendor:          CISCO-BIZLINK
           QSFP Part Number:     L45593-D218-D10
           QSFP Serial Number:   LCC2809G26F-A
           Device Type:          CX6-DX PSID(NAP0000000027)
           Firmware Version:     22.44.1700
           Part Number:          111-05341
           Hardware Revision:    20
           Serial Number:        032403001370

```

Step 5: Restore the storage system to normal operation

Restore your storage system to normal operation by giving back storage to the healthy controller, restoring automatic giveback, and reenabling AutoSupport automatic case creation.

Steps

1. Return the healthy controller (the controller that was taken over) to normal operation by giving back its storage:

```
storage failover giveback -ofnode healthy_node_name
```

2. Restore automatic giveback from the console of the impaired controller (the controller that took over the healthy controller):

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

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

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

Step 6: Return the failed part to NetApp

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

Replace an I/O module - AFF A20, AFF A30, and AFF A50

Replace an I/O module in your AFF A20, AFF A30, or AFF A50 storage system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

Before you begin

All other components in the storage system must be functioning properly; if not, contact [NetApp Support](#) before continuing with this procedure.

About this task

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

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

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

Step 1: Shut down the impaired controller

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 <code>y</code> 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: <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> The <i>-halt true</i> parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

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

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

```
metrocluster node show
```

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Step 2: Replace a failed I/O module


To replace a failed I/O module, locate it in the controller and follow the specific sequence of steps.

Steps

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

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

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



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the replacement I/O module into the target slot:
 - a. Align the I/O module with the edges of the slot.
 - b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O Module.

c. Turn the thumbscrew clockwise to tighten.

5. Cable the I/O module.

Step 3: Reboot the controller

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

Steps

1. Reboot the controller from the LOADER prompt:

```
bye
```



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

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

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

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

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

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

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

Step 4: Return the failed part to NetApp

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

Replace the NV battery - AFF A20, AFF A30, and AFF A50

Replace the NV battery in your AFF A20, AFF A30, or AFF A50 storage system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

Before you begin

All other components in the storage system must be functioning properly; if not, contact [NetApp Support](#) before continuing with this procedure.

About this task

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

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

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

Step 1: Shut down the impaired controller

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 <code>y</code> 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: <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> The <i>-halt true</i> parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

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

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

```
metrocluster node show
```

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Step 2: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

Steps


1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.


If you are disconnecting a...	Then...
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.


6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace the NV battery

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

Steps

- 1. If you are not already grounded, properly ground yourself.
- 2. Locate the NV battery.
- 3. Remove the NV battery:



1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	1. Push in and hold the tab on the connector. 2. Pull the connector up and out of the socket. As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.

- 4. Install the replacement NV battery:
 - a. Remove the replacement battery from its package.
 - b. Plug the wiring connector into its socket.
 - c. Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
 - d. Place the NV battery into its compartment.


The NV battery should sit flush in its compartment.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



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

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.



Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
 - a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

- b. Rotate the controller handles up and lock in place with the tabs.



The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

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

If you are reconnecting a...	Then...
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

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

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

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

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

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

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

Step 5: Return the failed part to NetApp

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

Hot-swap a power supply - AFF A20, AFF A30, and AFF A50

Replace an AC or DC power supply unit (PSU) in your AFF A20, AFF A30, or AFF A50 storage system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cord, replacing the faulty PSU, and then reconnecting it to the power source.

About this task

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

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

- **IMPORTANT:** Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.
- Use the appropriate procedure for your type of PSU: AC or DC.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

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

Option 1: Hot-swap an AC PSU

To replace an AC PSU, complete the following steps.


Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the power cord from the PSU by opening the power cord retainer, and then unplug the power cord from the PSU.



PSUs do not have a power switch.

4. Remove the PSU:



1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	Pull the PSU out of the controller while using your other hand to support its weight.  The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.

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

A PSU will only properly engage with the internal connector and lock in place one way.



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

- c. Rotate the handle down, so it is out of the way of normal operations.
6. Reconnect the power cord to the PSU and secure the power cord with the power cord retainer.

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

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

Option 2: Hot-swap a DC PSU

To replace a DC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the PSU:



PSUs do not have a power switch.

- a. Unscrew the two thumb screws on the D-SUB DC power cord connector.

The illustration and table in step 4 shows the two thumb screws (item #1) and the D-SUB DC power cord connector (item #2).


- b. Unplug the cord from the PSU and set it aside.

4. Remove the PSU:

- a. Rotate the handle up, to its horizontal position, and then grasp it.
- b. With your thumb, press the terracotta tab to release the locking mechanism.
- c. Pull the PSU out of the controller while using your other hand to support its weight.



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



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

5. Insert the replacement PSU:

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

A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.



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

- c. Rotate the handle down, so it is out of the way of normal operations.

6. Reconnect the D-SUB DC power cord:

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

- a. Plug the D-SUB DC power cord connector into the PSU.
 - b. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

Replace the real-time clock battery - AFF A20, AFF A30, and AFF A50

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your AFF A20, AFF A30, or AFF A50 storage system to ensure that services and applications relying on accurate time synchronization remain operational.

Before you begin

All other components in the storage system must be functioning properly; if not, contact [NetApp Support](#) before continuing with this procedure.

About this task

- You can use this procedure with all versions of ONTAP supported by your storage system.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

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

Step 1: Shut down the impaired controller

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 <code>y</code> 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: <pre>storage failover takeover -ofnode impaired_node_name -halt true</pre> The <i>-halt true</i> parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

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

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

```
metrocluster node show
```

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Step 2: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact [NetApp Support](#) before continuing with this procedure.

Steps


1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.



If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact [NetApp Support](#) before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller
---	-----------------------------------

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:



Power supplies (PSUs) do not have a power switch.


If you are disconnecting a...	Then...
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

Keep track of where the cables were connected.

5. Remove the impaired controller:

The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.


Step 3: Replace the RTC battery

Remove the failed RTC battery and install the replacement RTC battery.

You must use an approved RTC battery.

Steps

- 1. Locate the RTC battery.
- 2. Remove the RTC battery:



1	Gently rotate the RTC battery at an angle away from its holder.
2	Lift the RTC battery out of its holder.


- 3. Install the replacement RTC battery:
 - a. Remove the replacement battery from the antistatic shipping bag.
 - b. Position the battery so that the plus sign on the battery faces out to correspond with the plus sign on the motherboard.
 - c. Insert the battery into the holder at an angle, and then push it into an upright position so it is fully seated in the holder.
 - d. Visually inspect the battery to make sure that it is completely seated in its holder and that the polarity is correct.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.



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

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.



Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
 - a. Firmly push on the handles until the controller meets the midplane and is fully seated.



Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

- b. Rotate the controller handles up and lock in place with the tabs.



The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

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

If you are reconnecting a...	Then...
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

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

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

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

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

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

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

Step 5: Reset the time and date on the controller



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

1. On the healthy controller, check the date and time:

```
cluster date show
```



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

2. On the impaired controller, at the LOADER prompt, check the time and date:

```
cluster date show
```

- a. If necessary, modify the date:

```
set date mm/dd/yyyy
```

- b. If necessary, set the time, in GMT:

```
set time hh:mm:ss
```

c. Confirm the date and time.

3. At the LOADER prompt, enter `bye` to reinitialize the I/O modules, other components, and let the controller reboot.

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.

Key specifications

Key specifications for AFF A20

The following are select specifications for the AFF A20 storage system in a single high availability pair. Visit NetApp Hardware Universe (HWU) for the complete specifications for this storage system.

AFF A20 specifications at a glance

- Platform Configuration: AFF A20 Single Chassis HA Pair
- Max Raw Capacity: 4.0392 PB
- Memory: 128.0000 GB
- Form Factor: 2U chassis with 2 HA controllers and 24 drive slots
- ONTAP Version: ONTAP: 9.16.1P2
- PCIe Expansion Slots: 8
- Minimum ONTAP Version: ONTAP 9.16.1RC1

Scaleout maximums

- Type: NAS; HA Pairs: 3; Raw Capacity: 12.1 PB / 10.8 PiB; Max Memory: 384 GB
- Type: SAN; HA Pairs: 3; Raw Capacity: 12.1 PB / 10.8 PiB; Max Memory: 384 GB
- Type: HA Pair; Raw Capacity: 4.0 PB / 3.6 PiB; Max Memory: 128.0000

I/O

Onboard I/O

No onboard I/O data.

Total I/O

- Protocol: Ethernet 100 Gbps; Ports: 12
- Protocol: Ethernet 25 Gbps; Ports: 32
- Protocol: Ethernet 10 Gbps; Ports: 24
- Protocol: FC 64 Gbps; Ports: 24

- Protocol: NVMe/FC 64 Gbps; Ports: 24
- Ports: 0
- Protocol: SAS 12 Gbps; Ports: 16

Management ports

- Protocol: Ethernet 1 Gbps; Ports: 2
- Protocol: RS-232 115 Kbps; Ports: 4
- Protocol: USB 600 Mbps; Ports: 2

Storage networking supported

- CIFS
- FC
- iSCSI
- NFS v3
- NFS v4.0
- NFS v4.1
- NFS v4.2
- NVMe/FC
- NVMe/TCP
- S3
- S3 with NAS
- SMB 2.0
- SMB 2.1
- SMB 2.x
- SMB 3.0
- SMB 3.1
- SMB 3.1.1

System environment specifications

- Typical Power: 2489 BTU/hr
- Worst-case Power: 3890 BTU/hr
- Weight: 60.7 lb, 27.5 kg
- Height: 2U
- Width: 19" IEC rack-compliant (17.6" 44.7 cm)
- Depth:
- Operating Temp/Altitude/Humidity: 10°C to 35°C (50°F to 95°F) at up to 3048m (10000 ft) elevation; 8% to 80% relative humidity, noncondensing
- Non-operating Temp/Humidity: -40°C to 70°C (-40°F to 158°F) up to 12192m (40000 ft); 10% to 95% relative humidity, noncondensing, in original container

- Acoustic Noise: Declared sound power (LwAd): 8.0; Sound pressure (LpAm) (bystander positions): 70.5 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 15598-1, EN60825-1, EN62368-1, IEC 62368-1, IEC60950-1, IS 13252(part 1)

High availability

- Ethernet based baseboard management controller (BMC) and ONTAP management interface
- Redundant hot-swappable controllers
- Redundant hot-swappable power supplies
- SAS in-band management over SAS connections for external shelves

Key specifications for AFF A30

The following are select specifications for the AFF A30 storage system in a single high availability pair. Visit NetApp Hardware Universe (HWU) for the complete specifications for this storage system.

AFF A30 specifications at a glance

- Platform Configuration: AFF A30 Single Chassis HA Pair
- Max Raw Capacity: 4.0392 PB
- Memory: 128.0000 GB
- Form Factor: 2U chassis with 2 HA controllers and 24 drive slots
- ONTAP Version: ONTAP: 9.16.1P2
- PCIe Expansion Slots: 8
- Minimum ONTAP Version: ONTAP 9.16.1RC1

Scaleout maximums

- Type: NAS; HA Pairs: 4; Raw Capacity: 16.2 PB / 14.3 PiB; Max Memory: 512 GB
- Type: SAN; HA Pairs: 4; Raw Capacity: 16.2 PB / 14.3 PiB; Max Memory: 512 GB
- Type: HA Pair; Raw Capacity: 4.0 PB / 3.6 PiB; Max Memory: 128.0000

I/O

Onboard I/O

No onboard I/O data.

Total I/O

- Protocol: Ethernet 100 Gbps; Ports: 16
- Protocol: Ethernet 25 Gbps; Ports: 24
- Protocol: Ethernet 10 Gbps; Ports: 24
- Protocol: FC 64 Gbps; Ports: 24
- Protocol: NVMe/FC 64 Gbps; Ports: 24
- Ports: 0
- Protocol: SAS 12 Gbps; Ports: 16

Management ports

- Protocol: Ethernet 1 Gbps; Ports: 2
- Protocol: RS-232 115 Kbps; Ports: 4
- Protocol: USB 600 Mbps; Ports: 2

Storage networking supported

- CIFS
- FC
- iSCSI
- NFS v3
- NFS v4.0
- NFS v4.1
- NFS v4.2
- NVMe/FC
- NVMe/TCP
- S3
- S3 with NAS
- SMB 2.0
- SMB 2.1
- SMB 2.x
- SMB 3.0
- SMB 3.1
- SMB 3.1.1

System environment specifications

- Typical Power: 2892 BTU/hr
- Worst-case Power: 4445 BTU/hr
- Weight: 60.7 lb, 27.5 kg
- Height: 2U
- Width: 19" IEC rack-compliant (17.6" 44.7 cm)
- Depth:
- Operating Temp/Altitude/Humidity: 10°C to 35°C (50°F to 95°F) at up to 3048m (10000 ft) elevation; 8% to 80% relative humidity, noncondensing
- Non-operating Temp/Humidity: -40°C to 70°C (-40°F to 158°F) up to 12192m (40000 ft); 10% to 95% relative humidity, noncondensing, in original container
- Acoustic Noise: Declared sound power (LwAd): 8.0; Sound pressure (LpAm) (bystander positions): 70.5 dB

Compliance

- Certifications EMC/EMI: AMCA, FCC, ICES, KC, Morocco, VCCI
- Certifications safety: BIS, CB, CSA, G_K_U-SoR, IRAM, NOM, NRCS, SONCAP, TBS
- Certifications Safety/EMC/EMI: EAC, UKRSEPRO
- Certifications Safety/EMC/EMI/RoHS: BSMI, CE DoC, UKCA DoC
- Standards EMC/EMI: BS-EN-55032, BS-EN55035, CISPR 32, EN55022, EN55024, EN55032, EN55035, EN61000-3-2, EN61000-3-3, FCC Part 15 Class A, ICES-003, KS C 9832, KS C 9835
- Standards Safety: ANSI/UL60950-1, ANSI/UL62368-1, BS-EN62368-1, CAN/CSA C22.2 No. 60950-1, CAN/CSA C22.2 No. 62368-1, CNS 15598-1, EN60825-1, EN62368-1, IEC 62368-1, IEC60950-1, IS 13252(part 1)

High availability

- Ethernet based baseboard management controller (BMC) and ONTAP management interface
- Redundant hot-swappable controllers
- Redundant hot-swappable power supplies
- SAS in-band management over SAS connections for external shelves

Key specifications for AFF A50

The following are select specifications for the AFF A50 storage system in a single high availability pair. Visit NetApp Hardware Universe (HWU) for the complete specifications for this storage system.

AFF A50 specifications at a glance

- Platform Configuration: AFF A50 Single Chassis HA Pair
- Max Raw Capacity: 4.0392 PB
- Memory: 256.0000 GB
- Form Factor: 2U chassis with 2 HA controllers and 24 drive slots

- ONTAP Version: ONTAP: 9.16.1P2
- PCIe Expansion Slots: 8
- Minimum ONTAP Version: ONTAP 9.16.1RC1

Scaleout maximums

- Type: NAS; HA Pairs: 12; Raw Capacity: 48.5 PB / 43.0 PiB; Max Memory: 3072 GB
- Type: SAN; HA Pairs: 6; Raw Capacity: 24.2 PB / 21.5 PiB; Max Memory: 1536 GB
- Type: HA Pair; Raw Capacity: 4.0 PB / 3.6 PiB; Max Memory: 256.0000

I/O

Onboard I/O

No onboard I/O data.

Total I/O

- Protocol: Ethernet 100 Gbps; Ports: 16
- Protocol: Ethernet 25 Gbps; Ports: 24
- Protocol: Ethernet 10 Gbps; Ports: 24
- Protocol: FC 64 Gbps; Ports: 24
- Protocol: NVMe/FC 64 Gbps; Ports: 24
- Ports: 0
- Protocol: SAS 12 Gbps; Ports: 16

Management ports

- Protocol: Ethernet 1 Gbps; Ports: 2
- Protocol: RS-232 115 Kbps; Ports: 4
- Protocol: USB 600 Mbps; Ports: 2

Storage networking supported

- CIFS
- FC
- iSCSI
- NFS v3
- NFS v4.0
- NFS v4.1
- NFS v4.2
- NVMe/FC
- NVMe/TCP
- S3

- S3 with NAS
- SMB 2.0
- SMB 2.1
- SMB 2.x
- SMB 3.0
- SMB 3.1
- SMB 3.1.1

System environment specifications

- Typical Power: 3076 BTU/hr
- Worst-case Power: 4860 BTU/hr
- Weight: 60.7 lb, 27.5 kg
- Height: 2U
- Width: 19" IEC rack-compliant (17.6" 44.7 cm)
- Depth:
- Operating Temp/Altitude/Humidity: 10°C to 35°C (50°F to 95°F) at up to 3048m (10000 ft) elevation; 8% to 80% relative humidity, noncondensing
- Non-operating Temp/Humidity: -40°C to 70°C (-40°F to 158°F) up to 12192m (40000 ft); 10% to 95% relative humidity, noncondensing, in original container
- Acoustic Noise: Declared sound power (LwAd): 8.0; Sound pressure (LpAm) (bystander positions): 70.5 dB

Compliance

- Certifications EMC/EMI: AMCA, FCC, ICES, KC, Morocco, VCCI
- Certifications safety: BIS, CB, CSA, G_K_U-SoR, IRAM, NOM, NRCS, SONCAP, TBS
- Certifications Safety/EMC/EMI: EAC, UKRSEPRO
- Certifications Safety/EMC/EMI/RoHS: BSMI, CE DoC, UKCA DoC
- Standards EMC/EMI: BS-EN-55032, BS-EN55035, CISPR 32, EN55022, EN55024, EN55032, EN55035, EN61000-3-2, EN61000-3-3, FCC Part 15 Class A, ICES-003, KS C 9832, KS C 9835
- Standards Safety: ANSI/UL60950-1, ANSI/UL62368-1, BS-EN62368-1, CAN/CSA C22.2 No. 60950-1, CAN/CSA C22.2 No. 62368-1, CNS 15598-1, EN60825-1, EN62368-1, IEC 62368-1, IEC60950-1, IS 13252(part 1)

High availability

- Ethernet based baseboard management controller (BMC) and ONTAP management interface
- Redundant hot-swappable controllers
- Redundant hot-swappable power supplies
- SAS in-band management over SAS connections for external shelves

Copyright information

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

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

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

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

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

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

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

Trademark information

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