



Maintain

Install and maintain

NetApp
September 25, 2024

Table of Contents

- Maintain 1
 - Maintain FAS70 and FAS90 hardware 1
 - Boot media 2
 - Controller 19
 - Replace a DIMM - FAS70 and FAS90 31
 - Replace a fan - FAS70 and FAS90 36
 - Replace the FlashCache module carrier or a caching module - FAS70 and FAS90 37
 - Replace NVRAM - FAS70 and FAS90 44
 - Replace the NV Battery - FAS70 and FAS90 52
 - I/O module 57
 - Replace a power supply - FAS70 and FAS90 70
 - Replace the real-time clock battery - FAS70 and FAS90 71
 - Replace System management module - FAS70 and FAS90 77

Maintain

Maintain FAS70 and FAS90 hardware

You might need to perform maintenance procedures on your hardware. Procedures specific to maintaining your FAS70 and FAS90 system components are in this section.

The procedures in this section assume that the FAS70 and FAS90 system has already been deployed as a storage node in the ONTAP environment.

System components

For the FAS70 and FAS90 storage system, you can perform maintenance procedures on the following components.

| | |
|---|--|
| Boot media | The boot media stores a primary and secondary set of ONTAP image files that the system uses when it boots. |
| 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. |
| Fan | A fan cools the controller. |
| FlashCache | Flash Cache speeds access to data through real-time intelligent caching of recently read user data and NetApp metadata. It's effective for random read-intensive workloads, including databases, email, and file services. |
| NVRAM | The NVRAM (Non-Volatile Random Access Memory) is a module that allows the controller to protect and save in-flight data if the system loses power. The system ID resides in the NVRAM module. When replaced, the controller assumes the new system ID from the replacement NVRAM module. |
| NV battery | The NV battery is responsible for providing power to the NVRAM module while data in-flight is being destaged to flash memory after a power loss. |
| I/O module | The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller. |
| Power supply | A power supply provides a redundant power source in a controller. |
| Real-time clock battery | A real-time clock battery preserves system date and time information if the power is off. |

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

Boot media

Boot media replacement workflow - FAS70 and FAS90

Follow these workflow steps to replace your boot media.

1 Review boot media replacement requirements

To replace the boot media, you must meet certain requirements.

2 Check onboard encryption keys

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

3 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.

4 Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive to the replacement boot media.

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 ONATP 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.

Boot media replace requirements - FAS70 and FAS90

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

- You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz`.
- You must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct controller:
 - The *impaired* controller is the controller on which you are performing maintenance.
 - The *healthy* controller is the HA partner of the impaired controller.

Check onboard encryption keys - FAS70 and FAS90

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

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

Check NVE or NSE

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

Verify security key-manager configuration

Steps

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




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

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

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

4. If the `Key Manager` type displays `onboard` and the `Restored` column displays anything other than `true`:
 - a. Enter the onboard security key-manager sync command: `security key-manager onboard sync`



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

5. If the `Key Manager` type displays `external` and the `Restored` column displays anything other than `true`:
 - a. Restore the external key management authentication keys to all nodes in the cluster: `security key-manager external restore`

If the command fails, contact NetApp Support at mysupport.netapp.com.
 - b. Verify that the `Restored` column displays `true` for all authentication keys: `security key-manager key query`
 - c. You can safely shut down the impaired controller.

Shut down the impaired controller - FAS70 and FAS90

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

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



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

3. Take the impaired controller to the LOADER prompt:

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

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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

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

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


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

Replace the boot media - FAS70 and FAS90

To replace the boot media, you must remove the System Management module from the back of the system, remove the impaired boot media, install the replacement boot media in the System Management module.

Step 1: Replace the boot media

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



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

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



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

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

- c. Push the locking button, rotate the boot media all the way down and then release the locking button.
5. Reinstall the System Management module.
 - a. Align the module with the edges of the enclosure slot opening.
 - b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
 6. Rotate the cable management tray up to the closed position.
 - a. Recable the System Management module.

Step 2: Transfer the ONTAP image to the boot media

The replacement boot media that you installed is without an ONTAP image, You can transfer the ONTAP image to the replacement boot media by downloading the appropriate ONTAP service image from the [NetApp Support Site](#) to a USB flash drive and then to the replacement boot media.

Before you begin

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

Steps

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



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

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

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

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

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

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

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

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

controller to boot to LOADER.

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

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



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

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



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

Boot the recovery image - FAS70 and FAS90

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

Steps

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

The image is downloaded from the USB flash drive.

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

| If your system is running... | Then... |
|------------------------------|--|
| ONTAP 9.16.0 or earlier | <p>a. On the impaired controller, press <code>Y</code> when you see <code>Do you want to restore the backup configuration now?</code></p> <p>b. On the impaired controller, press <code>Y</code> when prompted to overwrite <code>/etc/ssh/ssh_host_ecdsa_key</code>.</p> <p>c. On the healthy partner controller, set the impaired controller to advanced privilege level: <code>set -privilege advanced</code>.</p> <p>d. On the healthy partner controller, run the restore backup command: <code>system node restore-backup -node local -target -address impaired_node_IP_address</code>.</p> <p>NOTE: If you see any message other than a successful restore, contact NetApp Support.</p> <p>e. On the healthy partner controller, return the impaired controller to admin level: <code>set -privilege admin</code>.</p> <p>f. On the impaired controller, press <code>y</code> when you see <code>Was the restore backup procedure successful?</code>.</p> <p>g. On the impaired controller, press <code>y</code> when you see <code>...would you like to use this restored copy now?</code>.</p> <p>h. On the impaired controller, press <code>y</code> when prompted to reboot the impaired controller and press <code>ctrl-c</code> for the Boot Menu.</p> <p>i. If the system does not use encryption, select <i>Option 1 Normal Boot.</i>, otherwise go to Restore key managers.</p> <p>j. Connect the console cable to the partner controller.</p> <p>k. Give back the controller using the <code>storage failover giveback -fromnode local</code> command.</p> <p>l. Restore automatic giveback if you disabled it by using the <code>storage failover modify -node local -auto-giveback true</code> command.</p> <p>m. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the <code>system node autosupport invoke -node * -type all -message MAINT=END</code> command.</p> <p>NOTE: If the process fails, contact NetApp Support.</p> |

| If your system is running... | Then... |
|------------------------------|--|
| ONTAP 9.16.1 or later | <p>a. On the impaired controller, press <i>y</i> when prompted to restore the backup configuration.</p> <p>After restore procedure is successful, this message will be seen on the console - <code>syncflash_partner: Restore from partner complete.</code></p> <p>b. On the impaired controller, press <i>y</i> when prompted to confirm if the restore backup was successful.</p> <p>c. On the impaired controller, press <i>y</i> when prompted to use the restored configuration.</p> <p>d. On the impaired controller, press <i>y</i> when prompted to reboot the node.</p> <p>e. On the impaired controller, press <i>y</i> when prompted to reboot the impaired controller and press <i>ctrl-c</i> for the Boot Menu.</p> <p>f. If the system does not use encryption, select <i>Option 1 Normal Boot.</i>, otherwise go to Restore key managers.</p> <p>g. Connect the console cable to the partner controller.</p> <p>h. Give back the controller using the <i>storage failover giveback -fromnode local</i> command.</p> <p>i. Restore automatic giveback if you disabled it by using the <i>storage failover modify -node local -auto-giveback true</i> command.</p> <p>j. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the <i>system node autosupport invoke -node * -type all -message MAINT=END</i> command.</p> <p>NOTE: If the process fails, contact NetApp Support.</p> |

Restore encryption - FAS70 and FAS90

Restore encryption on the replacement boot media.

Step 1: Restore onboard key manager

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



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

Steps

1. Connect the console cable to the target controller.
2. Select one of the following options to restore the onboard key manager configuration from the ONATP boot menu.

Option 1: Systems with onboard key manager server configuration

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

Before you begin

You need the following information while restoring the OKM configuration:

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

Steps

1. From the ONTAP boot menu select option 10:

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

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

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



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

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

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

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

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


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

```
Successfully recovered keymanager secrets.
```

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



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

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

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

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

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

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

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

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

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

There are no entries matching your query.



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

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

Option 2: Systems with external key manager server configuration

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

Before you begin

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

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

Steps

1. Select Option 11 from the ONTAP boot menu.

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

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

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

You may also see these prompts instead:

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

3. Supply the information for each of these prompts:

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

Example

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

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

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

```
-----BEGIN RSA PRIVATE KEY-----
MIIEpQIBAAKCAQEAoUleaajEG6QC2h2Zih0jEaGVtQUexNeoCFwKPomSePmjDNtrU
MSB1SlX3VgCuElHk57XPdq6xSbYl b kIb4bAgLztHEmUDOkGmXYAkblQ=
-----END RSA PRIVATE KEY-----
```

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

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

Enter the IP address for the KMIP server: 10.10.10.10

Enter the port for the KMIP server [5696]:

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

4. The recovery process will complete:

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

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

```

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

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

```

Step 2: Complete the boot media replacement

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

1. Check the console output:

| If the console displays... | Then... |
|----------------------------|--|
| The login prompt | Go to Step 6. |
| Waiting for giveback... | a. Log into the partner controller. b. Confirm the target controller is ready for giveback with the <i>storage failover show</i> command. |

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

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

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

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



If the command fails, contact Customer Support.

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

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

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

Return the failed part to NetApp - FAS70 and FAS90

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 - FAS70 and FAS90

Follow these workflow steps to replace your controller module.

1

Review controller replacement requirements

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

2

Shut down the impaired controller

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

3

Replace the controller

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

4

Restore and verify the system configuration

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

5

Give back the controller

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

6

Complete controller replacement

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

Controller replace requirements - FAS70 and FAS90

Review the requirements for the controller replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section [Choosing the correct recovery procedure](#) to determine whether you should use this controller replacement procedure.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You must replace a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- Because the boot device is located on the System Management module that is installed in the back of the system, you do not need to move the boot device when replacing a controller module.
- It is important that you apply the commands in these steps on the correct systems:

- The *impaired* controller is the controller that is being replaced.
- The *replacement* controller is the new controller that is replacing the impaired controller.
- The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text log file.

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

Shut down the impaired controller - FAS70 and FAS90

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

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



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

3. Take the impaired controller to the LOADER prompt:

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

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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

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

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

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


Replace the controller module - FAS70 and FAS90

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

Step 1: Remove the controller module

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

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



| | |
|----------|---------------------|
| 1 | NVRAM status LED |
| 2 | NVRAM attention LED |




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

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

- The NVRAM status LED flashes when power is removed from the controller module and the controller module is in the "waiting for giveback" state, or the controller module is not taken over or halted properly (uncommitted data).
 - The NVRAM status LED flashes when the controller module is removed from the enclosure and could mean the controller module is not taken over or halted properly (uncommitted data). Confirm that the controller module has been cleanly takeover by the partner controller module or the impaired controller module shows `waiting for giveback`. Then, the flashing LED can be ignored (and the controller can be removed from the enclosure).
2. If you are not already grounded, properly ground yourself.
 3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

The controller module moves slightly out of the enclosure.



| | |
|----------|-----------------------|
| 1 | a Locking cam latches |
|----------|-----------------------|

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

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


Step 2: Move the fans

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

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



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



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

4. Install the fan in the replacement controller module:


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

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

Step 3: Move the NV battery

Move the NV battery to the replacement controller.

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




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

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

Step 4: Move system DIMMs

Move the DIMMs to the replacement controller module.

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



| | |
|----------|-------------|
| 1 | System DIMM |
|----------|-------------|

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



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

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

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



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

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

Step 5: Install the controller module

Reinstall the controller module and boot it.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.

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



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

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

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

Restore and verify the system configuration - FAS70 and FAS90

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

Step 1: Set and verify system time

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

About this task

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

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

Steps

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

The date and time are based on the configured timezone.

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

The date and time are given in GMT.

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

The date and time are given in GMT.

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

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

1. Boot to Maintenance mode on the replacement controller module and verify that all components display the same HA state: `ha-config show`

The HA state should be the same for all components.

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

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

- `ha`
 - `mcc` (not supported)
 - `mccip` (not supported in ASA systems)
 - `non-ha` (not supported)
3. If the displayed system state of the controller module does not match your system configuration, set the HA state for the controller module: `ha-config modify controller ha-state`
 4. Confirm that the setting has changed: `ha-config show`

Give back the controller - FAS70 and FAS90

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

Steps

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

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



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

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

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

3. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

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

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

This is required if both of the following are true:

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

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

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

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

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

4 entries were displayed.
```

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

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

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

Complete controller replacement - FAS70 and FAS90

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

Step 1: Verify LIFs and and check cluster health

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, check the cluster health, and reset automatic giveback.

Steps

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

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

2. Check the health of your cluster. See the [How to perform a cluster health check with a script in ONTAP KB](#) article for more information.
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 2: Return the failed part to NetApp

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

Replace a DIMM - FAS70 and FAS90

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

Before you begin

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

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

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



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

3. Take the impaired controller to the LOADER prompt:

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

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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


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

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

Step 2: Remove the controller module

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

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



| | |
|----------|---------------------|
| 1 | NVRAM status LED |
| 2 | NVRAM attention LED |




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

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

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

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

The controller module moves slightly out of the enclosure.



1

a Locking cam latches

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

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


Step 3: Replace a DIMM

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

- If you are not already grounded, properly ground yourself.
- Open the controller air duct on the top of the controller.
 - Insert your fingers in the recesses at the far ends of the air duct.
 - Lift the air duct and rotate it upward as far as it will go.
- Locate the DIMMs on your controller module and identify the DIMM for replacement.
- Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.



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



1

DIMM and DIMM ejector tabs

- Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

The notch among the pins on the DIMM should line up with the tab in the socket.

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

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



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

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Close the controller air duct.

Step 4: Install the controller

Reinstall the controller module and boot it.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans



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

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

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

Step 5: Return the failed part to NetApp

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

Replace a fan - FAS70 and FAS90

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

Steps

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

Facing the controller module, fan modules are numbered 1 through 5, from left to right.




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

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



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



| | |
|--|----------------------|
| | Black release button |
|--|----------------------|

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

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

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

Replace the FlashCache module carrier or a caching module - FAS70 and FAS90

The NVMe SSD FlashCache module carrier contains one or two FlashCache modules (caching modules) with a single SSD FlashCache drive integrated into each caching module.

The FAS70 supports 2TB caching modules and FAS90 supports 4TB caching modules. You cannot mix caching modules of different capacity in the FlashCache module carrier.

You can perform either of the following procedures depending on what needs to be replaced: the entire Flashcache module carrier or a caching module.

- [Replace the FlashCache module carrier](#)
- [Replace the caching module](#)

Replace the FlashCache module carrier

The FlashCache module carrier is located in slot 6 and houses up to two FlashCache modules. You cannot hot-swap the FlashCache module carrier

Before you begin

- Ensure your storage system has the appropriate operating system for the replacement FlashCache module carrier.
- Confirm all other components are functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired node

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

Option 1 Most systems

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

About this task

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

Steps

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

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

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

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

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

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

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

Option 2 MetroCluster configuration



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

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

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

Steps

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

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

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

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

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

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


Step 2: Replace the FlashCache module carrier

Perform the following steps to replace the FlashCache module carrier.

Steps

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

2. Locate the failed FlashCache module carrier, in slot 6, by the lit amber Attention LED on the front of the FlashCache module carrier.



| | |
|----------|--------------------------------------|
| 1 | FlashCache module carrier |
| 2 | Caching module slot numbers |
| 3 | FlashCache module carrier cam handle |
| 4 | FlashCache module carrier fault LED |

3. Remove the failed FlashCache module carrier:
 - a. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
 - b. Pinch the blue tab at the bottom of the FlashCache module carrier.
 - c. Rotate the tab away from the module.
4. Pull the FlashCache module carrier out of the controller module and set it on an antistatic mat.
5. Move the caching modules to the replacement FlashCache module carrier:
 - a. Pinch the Terra Cotta tab at the top of the caching module and rotate the cam handle away from the caching module.
 - b. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the FlashCache module carrier.
 - c. Install the caching module into the same slot in the replacement FlashCache module carrier and rotate the cam handle to the closed position on the caching module to lock it in place.
6. Repeat these steps if there is a second caching module.
7. Install the replacement FlashCache module carrier into the system:
 - a. Align the module with the edges of the enclosure slot opening.
 - b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the

way up to lock the module in place.

- c. Rotate the cable management tray up to the closed position.

Step 3: Reboot the controller

After you replace the FlashCache module carrier, you must reboot the controller module.

Steps

1. From the LOADER prompt, reboot the node: *bye*



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

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

Step 4: Return the failed part to NetApp

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

Replace the caching module

Before you begin

The FlashCache modules (caching modules) are located in slot 6-1 or in slot 6-2 or in both slot 6-1 and slot 6-2.

You can hot-swap the individual caching modules with caching modules of the same capacity from the same vendor or from a different supported vendor.

Before you begin

- Ensure the replacement caching module has the same capacity as the failed one, from the same vendor or from a different supported vendor.
- Confirm all other components are functioning properly; if not, you must contact technical support.
- The drives in the caching modules are not field replaceable units (FRU). You must replace the entire caching module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the failed caching module, in slot 6, by the lit amber Attention LED on the front of the caching module.
3. Prepare the caching module slot for replacement as follows:
 - a. Record the caching module capacity, part number, and serial number on the target node: *system node run local sysconfig -av 6*
 - b. In admin privilege level, prepare the target caching module slot for removal, responding *y* when prompted whether to continue: *system controller slot module remove -node node_name -slot slot_number* The following command prepares slot 6-1 on node1 for removal, and displays a message that it is safe to remove:

```
::> system controller slot module remove -node node1 -slot 6-1
```

Warning: SSD module in slot 6-1 of the node node1 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.


- c. Display the slot status with the `system controller slot module show` command.

The caching module slot status displays `powered-off` in the screen output for the caching module that needs replacing.



See the [Command man pages](#) for your version of ONTAP for more details.

4. Remove the caching module:



| | |
|----------|---------------------------|
| 1 | Caching module cam handle |
| 2 | Caching module fault LED |

- Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
- Press the terra cotta release button on the front of the caching module.
- Rotate the cam handle as far as it will go.
- Remove the caching module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the FlashCache module carrier.

Be sure to support the caching module as you remove it from the FlashCache module carrier.

5. Install the replacement caching module:

- a. Align the edges of the caching module with the opening in the controller module.
 - b. Gently push the caching module into the bay until the cam handle engages.
 - c. Rotate the cam handle until it locks into place.
 - d. Rotate the cable management tray up to the closed position.
6. Bring the replacement caching module online by using the `system controller slot module insert` command as follows:

The following command prepares slot 6-1 on node1 for power-on, and displays a message that it is powered on:

```
::> system controller slot module insert -node node1 -slot 6-1

Warning: NVMe module in slot 6-1 of the node localhost 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.
```

7. Verify the slot status using the `system controller slot module show` command.

Make sure that command output reports status for the as `powered-on` and ready for operation.

8. Verify that the replacement caching module is online and recognized, and then visually confirm that the amber attention LED is not lit: `sysconfig -av slot_number`



If you replace the caching module with a caching module from a different vendor, the new vendor name is displayed in the command output.

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

Replace NVRAM - FAS70 and FAS90

The NVRAM module consists of the NVRAM12 hardware and field-replaceable DIMMs. You can replace a failed NVRAM module or the DIMMs inside the NVRAM module. To replace a failed NVRAM module, you must remove the module from the enclosure, move the DIMMs to the replacement module, and install the replacement NVRAM module into the enclosure.

All other components in the system must be functioning properly; if not, you must contact [NetApp Support](#).

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

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



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

3. Take the impaired controller to the LOADER prompt:

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

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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

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

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


Step 2: Replace the NVRAM module

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

1. If you are not already grounded, properly ground yourself.
2. Unplug the power cord from both PSUs.
3. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the tray down.
4. Remove the impaired NVRAM module from the enclosure:
 - a. Depress the locking cam button.

The cam button moves away from the enclosure.

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



| | |
|----------|--------------------|
| 1 | Cam locking button |
| 2 | DIMM locking tabs |

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


Step 3: Replace a NVRAM DIMM

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

1. If you are not already grounded, properly ground yourself.
2. Unplug the power cord from both PSUs.
3. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the

tray down.

4. Remove the target NVRAM module from the enclosure.



| | |
|----------|--------------------|
| 1 | Cam locking button |
| 2 | DIMM locking tabs |

5. Set the NVRAM module on a stable surface.
6. Locate the DIMM to be replaced inside the NVRAM module.



Consult the FRU map label on the side of the NVRAM module to determine the locations of DIMM slots 1 and 2.

7. Remove the DIMM by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.
8. Install the replacement DIMM by aligning the DIMM with the socket and gently pushing the DIMM into the socket until the locking tabs lock in place.
9. Install the NVRAM module into the enclosure:
 - a. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.
10. Recable the PSUs.
11. Rotate the cable management tray up to the closed position.

Step 4: Reboot the controller

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

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

Step 5: Reassign disks

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



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

Steps

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

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

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

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

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

The controller takes back its storage and completes booting.

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



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

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

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

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

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

The disks belonging to the controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 151759706:

```
node1:> storage disk show -ownership

Disk Aggregate Home Owner DR Home Home ID Owner ID DR Home ID
Reserver Pool
-----
-----
-----
-----
-----
1.0.0 aggr0_1 node1 node1 - 151759706 151759706 -
151759706 Pool0
1.0.1 aggr0_1 node1 node1 151759706 151759706 -
151759706 Pool0
.
.
.
```

6. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

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

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

This is required if both of the following are true:

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

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

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

```

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

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

4 entries were displayed.

```

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

Step 6: Return the failed part to NetApp

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

Replace the NV Battery - FAS70 and FAS90

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

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

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



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

3. Take the impaired controller to the LOADER prompt:

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

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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


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

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

Step 2: Remove the controller module

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

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



| | |
|----------|---------------------|
| 1 | NVRAM status LED |
| 2 | NVRAM attention LED |




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

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

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

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

The controller module moves slightly out of the enclosure.



| | |
|----------|-----------------------|
| 1 | a Locking cam latches |
|----------|-----------------------|


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

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

Step 3: Replace the NV battery

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

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



| | |
|----------|---------------------------|
| 1 | NV battery air duct cover |
| 2 | NV battery plug |

2. Lift the battery up to access the battery plug.

3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.

4. Lift the battery out of the air duct and controller module, and then set it aside.

5. Remove the replacement battery from its package.

6. Install the replacement battery pack into the controller:

a. Plug the battery plug into the riser socket and make sure that the plug locks into place.

b. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.

7. Close the NV air duct cover.

Make sure that the plug locks into the socket.

Step 4: Reinstall the controller module

Reinstall the controller module and boot it.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans



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

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

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

Step 5: Return the failed part to NetApp

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

I/O module

Overview of add and replace I/O module - FAS70 and FAS90

You can replace a failed I/O module in your storage system with the same type of I/O module, or with a different kind of I/O module. You can also add an I/O module into a system with empty slots.

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

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

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

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

Add I/O module - FAS70 and FAS90

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

Before you begin

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

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

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

Step 1: Shut down the impaired controller module

Shut down or take over the impaired controller module.

Option 1: Most systems

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

About this task

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

Steps

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

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

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

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

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

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

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

Option 2: Controller is in a MetroCluster



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

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

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

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

Steps

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

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

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

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

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

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


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

Step 2: Add I/O modules

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



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



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

- a. Depress the cam latch on the blanking module in the target slot.
 - b. Rotate the cam latch away from the module as far as it will go.
 - c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
4. Install the I/O module:
- a. Align the I/O module with the edges of the enclosure slot opening.
 - b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
5. Cable the I/O module.

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

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



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

6. Rotate the cable management tray up to the closed position.
7. Reboot the controller from the LOADER prompt: `bye`



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

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

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

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

1. If you are:

| Replacing a... | Then... |
|---|---|
| NIC I/O module with the same the same number of ports | The LIFs will automatically migrate when its controller module is shut down. |
| NIC I/O module with fewer ports | Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for information about using System Manager to permanently move the LIFs. |
| NIC I/O module with a storage I/O module | Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF . |

Step 1: Shut down the impaired controller module

Shut down or take over the impaired controller module.

Option 1: Most systems

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

About this task

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

Steps

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

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

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

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

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

Option 2: Controller is in a MetroCluster



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

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

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

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

Steps

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

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

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

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

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

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


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

Step 2: Replace an I/O module

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



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



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

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

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

5. Install the I/O module into the target slot in the enclosure:
 - a. Align the module with the edges of the enclosure slot opening.
 - b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
6. Cable the I/O module.
7. Repeat the remove and install steps to replace additional modules for the controller.
8. Rotate the cable management tray up to the closed position.
9. Reboot the controller from the LOADER prompt: `_bye_`
 - a. Check the version of BMC on the controller: `system service-processor show`
 - b. Update the BMC firmware if needed: `system service-processor image update`
 - c. Reboot the node: `bye`



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



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

10. Give back the controller from the partner controller: `storage failover giveback -ofnode target_node_name`
11. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto-giveback true`
12. If you added:

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

13. Repeat these steps for controller B.

Replace I/O module - FAS70 and FAS90

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

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

Step 1: Shut down the impaired node

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

Option 1: Most systems

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

About this task

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

Steps

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

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

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

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

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

Option 2: Controller is in a MetroCluster



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

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

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

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

Steps

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

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

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

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

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

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

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


Step 2: Replace a failed I/O module

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

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



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



| | |
|----------|---------------|
| 1 | I/O cam latch |
|----------|---------------|

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

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

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

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

Step 3: Reboot the controller

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



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

Steps

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



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

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

Step 4: Return the failed part to NetApp

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

Replace a power supply - FAS70 and FAS90

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

The power supplies are redundant and hot-swappable.

About this task

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




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

Steps

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



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



1

Terracotta PSU locking tab

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

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



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

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

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

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

Replace the real-time clock battery - FAS70 and FAS90

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

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

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



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

3. Take the impaired controller to the LOADER prompt:

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

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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


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

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

Step 2: Remove the controller module

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

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



| | |
|----------|---------------------|
| 1 | NVRAM status LED |
| 2 | NVRAM attention LED |




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

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

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

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

The controller module moves slightly out of the enclosure.



| | |
|----------|-----------------------|
| 1 | a Locking cam latches |
|----------|-----------------------|


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

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

Step 3: Replace the RTC battery

Remove failed RTC battery and install the replacement RTC battery.

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



| | |
|----------|-------------------------|
| 1 | RTC battery and housing |
|----------|-------------------------|

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



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

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

Step 4: Reinstall the controller module

Reinstall the controller module and boot it.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans



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

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

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

Step 5: Reset the time and date on the controller



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

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



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

- a. At the LOADER prompt on the target controller, check the time and date with the `cluster date show` command.
- b. If necessary, modify the date with the `set date mm/dd/yyyy` command.
- c. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
 1. Confirm the date and time on the target controller.
 2. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
 3. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode _impaired_node_name_`
 4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 6: Return the failed part to NetApp

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

Replace System management module - FAS70 and FAS90

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

The System Management module has the following onboard components:

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

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

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

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

Before you begin

- This procedure uses the following terminology:
 - The impaired controller is the controller on which you are performing maintenance.
 - The healthy controller is the HA partner of the impaired controller.
- All other system components must be working properly.
- The partner controller must be able to take over the impaired controller.
- You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

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



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

3. Take the impaired controller to the LOADER prompt:

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

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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

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

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


Step 2: Replace the impaired System Management module

Replace the impaired system management module.

1. Remove the System Management module:



Make sure NVRAM destage has completed before proceeding.



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


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



Make sure NVRAM destage has completed before proceeding.

- b. Remove any cables connected to the System Management module. Make sure that label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
- c. Disconnect the power cords from the PSU for the impaired controller.
- d. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
- e. Depress the cam button on the System Management module.
- f. Rotate the cam lever down as far as it will go.
- g. Loop your finger into the hole on the cam lever and pull the module straight out of the system.
- h. Place the System Management module on an anti-static mat, so that the boot media is accessible.

2. Move the boot media to the replacement System Management module:



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

| | |
|----------|---------------------------|
| 2 | Boot media locking button |
| 3 | Boot media |

- a. Press the blue boot media locking button in the impaired System Management module.
 - b. Rotate the boot media up and slide it out of the socket.
3. Install the boot media in the replacement System Management module:
- a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
 - b. Rotate the boot media down until it touches the locking button.
 - c. Depress the blue locking and rotate the boot media all the way down and release the blue locking button.
4. Install the replacement System Management module into the enclosure:
- a. Align the edges of the replacement System Management module with the system opening and gently push it into the controller module.
 - b. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.
5. Rotate the cable management arm up to the closed position.
6. Recable the System Management module.

Step 3: Reboot the controller module

Reboot the controller module.

1. Plug the power cables back into the PSU.

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

Step 4: Install licenses and register serial number

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

About this task

Until you install license keys, features requiring standard licenses continue to be available to the node.

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

Before you begin

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

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

Steps

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



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

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

Step 5: Return the failed part to NetApp

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

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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