



# Upgrading a system in an HA pair to a different system in an HA pair by moving disk shelves

You can upgrade a storage system in an HA pair to a storage system in an HA pair in another family by replacing the controllers.

## Before you begin

You should have collected system performance data by running the Performance and Statistics Collector (Perfstat) tool during normal peak time; you must have performed this task one to three weeks before upgrading the system.

Perfstat collects performance information and writes it to a text file, which can help troubleshoot any issues after the upgrade. You can find information about Perfstat and download the tool from the Downloads section of the the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com).

## About this task

This procedure is intended for systems running Data ONTAP 8.2.x operating in 7-Mode. FAS8080 systems require Data ONTAP 8.2.2 or later.

This procedure is written with the following assumptions:

- You are upgrading a system without internal storage and with disk shelves supported by the new system.  
If you want to upgrade a system with internal storage or with disk shelves not supported by the new system, see *Upgrading a FAS2240 system in an HA Pair with internal storage to another system in an HA pair by copying volumes*.
- You are upgrading to a system with the following characteristics:
  - It is new and unused.
  - It supports at least as many disks as the original system.  
The *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com) lists the number of disks supported by each system.
  - It might have internal storage or attached disk shelves.
  - It runs Data ONTAP 8.2 or later.
  - It will have the same name and IP address as the original system.
- Your system might be in an FC or iSCSI environment.  
If your system is in an FC or iSCSI environment, you need to take specific steps included in this procedure, and you must use the root volume of the original system on the upgraded system. See the *Data ONTAP SAN Configuration Guide for 7-Mode* and the *Data ONTAP SAN Administration Guide for 7-Mode* for more information about setting up and managing systems in SAN environments.
- Your current system and the system to which you are upgrading might have Storage Encryption on their disk drives.
- You have basic UNIX administration skills and storage knowledge.

**Note:** This document uses the term *HA pair* to describe two-controller storage systems in which one controller takes over and serves data if the other controller fails. Before the release of Data ONTAP 8.0, such systems were called *active/active configurations*.

## Steps

1. [Guidelines for upgrading the controllers](#) on page 2
2. [Gathering tools and documentation](#) on page 3
3. [Preparing the site and equipment for the upgrade](#) on page 4

#### 4. [Replacing the controller](#) on page 9

## Guidelines for upgrading the controllers

To upgrade the original nodes, you must follow certain guidelines and be aware of restrictions that affect the procedure.

### Supported upgrade paths

You can upgrade systems as follows:

From	To
FAS22xx	FAS25xx, 3220, 3250, 6220, 6250, 6290, FAS80xx
FAS31xx	3220, 3250, 6220, 6250, 6290, FAS80xx
FAS3210	3220, 3250, 6220, 6250, 6290, FAS80xx
3220	3250, 6220, 6250, 6290, FAS80xx
FAS3240	FAS3250
3240, 3270	3250, 6220, 6250, 6290, FAS80xx
3250	6220, 6250, 6290, FAS80xx
6040, 6080	6220, 6250, 6290, FAS80xx
FAS6210	6220, FAS8060, FAS8080
6220	FAS8060, FAS8080
FAS6240	6250, 6290, FAS8080
FAS6280	FAS6290
6280	FAS8080

**Note:** The system you are upgrading to should be new and unused.

**Note:** If the new system has fewer slots than the original system, or if it has fewer or different ports, you might need to add an adapter to the new system. See the *Hardware Universe* on the NetApp Support Site for details about specific platforms.

### Licensing in Data ONTAP 8.2.x

Starting with Data ONTAP 8.2, all license keys are 28 uppercase alphabetic characters in length.

In Data ONTAP 8.2.x, you need to obtain a new, 28-character license key for each package. You can get license keys from the NetApp Support Site in the *My Support* section under *Software licenses*. If the site does not have the license keys you need, contact your NetApp sales representative.

You have 90 days to install the matching keys for each controller. During the grace period, the new controller can use the same licensed functionality that the original controller had. After 90 days, the original controller's keys will be removed.

During the grace period, you have 24 hours after the first valid license key is installed to complete the license installation for all the packages that you want the new controller to use. After the 24-hour period, all previously installed licenses that were associated with the original system serial number will be removed.

You can use the `license show` command to check the time available before the grace period expires.

**Note:** For systems in HA pair, you need to install licenses for both controllers.

For detailed information about licensing in Data ONTAP 8.2, see the *Data ONTAP System Administration Guide for 7-Mode* and the *Data ONTAP Release Notes for 7-Mode*.

## Storage Encryption

The original nodes or the new nodes might be enabled for Storage Encryption. In that case, you need to take additional steps in this procedure to ensure that Storage Encryption is set up properly.

If you want to use Storage Encryption, all the disk drives associated with the nodes must have self-encrypting disk drives.

## Limitations

If your system is in an ESX environment and you upgrade the original system using the new system's root volume, the LUN serial numbers will change after the upgrade to reflect a serial number based on the new NVRAM ID. This causes an unnecessary resignature event that should be avoided. See the article *What causes ESX VMFS to "disappear"?* on the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com).

# Gathering tools and documentation

Before beginning the upgrade process, you need to gather the necessary tools and recommended documentation.

## Steps

1. Gather the tools you need to perform the upgrade:
  - Grounding strap
  - #2 Phillips screwdriver
2. Download from the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com) the documents that contain information helpful during the upgrade.

Download the version of the document for the version of Data ONTAP that the system is running.

Document	Contents
<i>Site Requirements Guide</i>	Contains information about the physical requirements of storage systems.
<i>Data ONTAP Storage Management Guide for 7-Mode</i>	Contains instructions for stopping deduplication processes and reassigning disks. Also contains information about Storage Encryption.
<i>Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode</i>	Contains information about quiescing SnapMirror.
<i>Data ONTAP Software Setup Guide for 7-Mode</i>	Contains information about configuring the Service Processor (SP) and about Storage Encryption.
<i>Data ONTAP System Administration Guide for 7-Mode</i>	Contains instructions for accessing the messages log, obtaining information to set up the SP, editing files in the <code>/etc</code> directory, and for performing other administrative tasks.
<i>Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode</i>	Contains instructions for downloading Data ONTAP.
<i>Installation and Setup Instructions</i> for the new system.	Contains platform-specific instructions for installing and cabling the new system. Included in the box with the controller and also available on the the NetApp Support Site.
The appropriate disk shelf guide	Contains instructions for installing and monitoring disk shelves and replacing disk shelf devices.

Document	Contents
<i>Universal SAS and ACP Cabling Guide</i>	Contains information for cabling SAS shelves applicable to all platforms.
<i>Data ONTAP SAN Administration Guide for 7-Mode</i>	Contains instructions for configuring and managing iSCSI and FC protocols for SAN environments.
<i>Data ONTAP SAN Configuration Guide for 7-Mode</i>	Contains information about FC and iSCSI topologies and wiring schemes.
<i>Data ONTAP 7-Mode High-Availability Configuration Guide</i>	Contains cabling instructions and other information for HA pairs.

The NetApp Support Site also contains documentation about disk shelves, NICs, and other hardware that you might use with your system. It also contains the *Hardware Universe*, which provides information about the hardware that the new system supports.

## Preparing the site and equipment for the upgrade

Before you upgrade the system, you need to ensure that you provide sufficient space for the new controllers. You also need to obtain information for setting up the remote management devices, and you might need to update Data ONTAP on the original system.

### Before you begin

You must have performed certain steps before starting the procedure:

- **Planning port migration**  
Because the ports on the original system might not match the ports on the new system, you must have planned port migration from the original system to the new system. During the procedure, you match the port locations on the original system to those on the new system.  
If the original system has more onboard GbE ports per controller than the new system, and if all the onboard GbE ports on the original system are being used, you must have added one or more GbE adapters to the new system.
- **Verifying MetroCluster configuration**  
If you are upgrading a system in a MetroCluster configuration, you need to verify that the switches are supported by the new controllers and the version of Data ONTAP that the controllers run. You also need to update the firmware on switches and any ATTO FibreBridge 6500N bridges.

### About this task

This procedure assumes that the original system is running an older version of Data ONTAP than the new system and includes steps for upgrading Data ONTAP on the original system, if necessary.

However, if the original system is running the desired version of Data ONTAP and the new system is running an older version, you need to update Data ONTAP on the new system before upgrading the controller hardware. See the *Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode* for information.

### Steps

1. Ensure that all cables connected to the back of the existing controllers are clearly labeled with the controller and port names.
2. Perform the following substeps to check the status of the HA pair:
  - a. Check if either controller is in takeover mode by entering the following command:

```
cf status
```

In takeover mode, a controller is actively taking over for its partner and displays a message similar to the following example:

```
node(takeover)>
local_node has taken over partner_node
```

Additional lines in the output might give the reason for the failover or indicate that the partner node is ready for giveback.

- b. If one of the controllers is in takeover mode, perform giveback by entering the following command:

```
cf giveback
```

- c. Make sure that both controllers are up by entering the following command on both controllers:

```
cf status
```

When both controllers are up and neither is in takeover mode, the system displays the following message:

```
local_node> cf status
Controller Failover enabled, partner_node is up
VIA Interconnect is up (link 0 up, link 1 up).
```

**Note:** Depending on the storage system model, the output might display either `RDMA interconnect` or `VIA interconnect` in the last line.

3. Ensure that all deduplication processes are stopped.

See the *Data ONTAP Storage Management Guide for 7-Mode* for information about deduplication.

4. Ensure that all SnapMirror processes are quiesced.

See the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode* for information about SnapMirror.

5. Review the messages log to make sure that there are no critical system issues.

See the *Data ONTAP System Administration Guide for 7-Mode* for information about accessing the message log.

6. Send an AutoSupport message to NetApp by entering the following command:

```
options autosupport.doit starting_disruptive_controller_head_upgrade
```

7. Obtain system information and backup configuration by entering the following commands and recording their output:

```
disk show -v
vol status
storage show disk -p
options
sysconfig -a
sysconfig -c
sysconfig -r
aggr status
ifconfig -a
system node service-processor show|rlm status
config dump -v config_filename
```

The file containing system information and the backup configuration saved with the `config dump` command is saved in the `/mroot/etc/configs` directory. You can give the file any name you choose.

8. Measure the amount of space that you have in your equipment rack or system cabinet.

You must allow for the appropriate amount of space for the new system.

9. Check the *Hardware Universe* on the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com) to ensure that the new system supports the disk shelves and any other hardware that you are migrating from the original system.
10. If you want to migrate a tape device, verify support for the new configuration on the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com) and with the backup application vendor.
11. Obtain IP addresses, mailhost addresses, and other information for the SPs on the system.
 

**Note:** You might want to reuse the network parameters of the RLMs or SPs from the original system for the SP on the new system.
12. Check the system invoice for the version of Data ONTAP installed on the new system.
13. Take one of the following actions:

---

**If the version of Data ONTAP on the original system is... Then...**

---

The same as the version on the new system    Go to Step 14.

---

Different from the version on the new system    Go to the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com) and download the version of Data ONTAP for the original system that matches the version installed on the new system.

**Note:** Downloads of the same version of Data ONTAP differ according to system model. You must choose the correct file for the model of your system.

**Note:** Make sure that you download the software image package to a directory that is accessible to the original system. See the *Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode* for instructions about how to download Data ONTAP.

---

14. If you want the new controllers to have the same licensed functionality as the original controllers, enter the following command to display a list of licenses and capture its output:

**license show**

15. Obtain new license keys for the new controllers from the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com).

If the site does not have the license keys you need, contact your NetApp sales representative.

16. Back up the data on the original system.
17. Save any information about site-specific customizations that you might want to reuse or save for reference, including NAS shares or SAN identity information, by taking one of the following actions:

---

**If the system... Then...**

---

Is not in a SAN environment

Enter the following commands and note their output:

**rdfile /etc/rc**

**rdfile /etc/hosts**

**rdfile /etc/exports**

**cifs shares**

---

If the system...	Then...
Is in a SAN environment	Enter the following commands and note their output: <pre>rdfile /etc/rc rdfile /etc/hosts rdfile /etc/exports cifs shares fcadmin config fcp nodename fcp portname show -v igroup show -v iscsi nodename lun show -v ucadmin show</pre>

18. Take one of the following actions:

If the original system...	Then...
Uses Storage Encryption	Complete Step 19 through Step 27, as appropriate, on both controllers of the original system. Complete all the steps first on one controller, and then all the steps on the other controller.  <b>Note:</b> Contact technical support to perform an optional step to preserve the security of the encrypted drives by rekeying all drives to a known authentication key.
Does not use Storage Encryption	Complete Step 23 through Step 27, as appropriate, on both controllers of the original system. Complete all the steps first on one controller, and then all the steps on the other controller.

19. Display the key ID for each self-encrypting disk on the original system by entering the following command at the system prompt:

```
disk encrypt show
```

**Example**

```
disk encrypt show
Disk      Key ID                                     Locked?
0c.00.1   0x0                                         No
0c.00.0   080CF0C80000000001000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
0c.00.3   080CF0C80000000001000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
0c.00.4   080CF0C80000000001000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
0c.00.2   080CF0C80000000001000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
0c.00.5   080CF0C80000000001000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
```

The first disk in the example is associated with an MSID (the default Security ID set by the manufacturer); the others are associated with a non-MSID. When a disk is locked, it requires authentication at the next disk power-on or power-cycle event. Both disks associated with an MSID and disks associated with a non-MSID can be locked or unlocked.

20. Examine the output of the `disk encrypt show` command, and if any disks are associated with a non-MSID key, rekey them to an MSID key by taking one of the following actions:

- To rekey disks individually, enter the following command, once for each disk:  

```
disk encrypt rekey 0x0 disk_name
```
- To rekey all disks at once, enter the following command:  

```
disk encrypt rekey 0x0 *
```

- Ensure that all the self-encrypting disks are associated with an MSID by entering the following command and examining the output:

```
disk encrypt show
```

### Example

The following example shows the output of the `disk encrypt show` command when all self-encrypting disks are associated with an MSID:

```
cluster::> disk encrypt show
Disk      Key ID                                     Locked?
-----
0b.10.23  0x0                                       No
0b.10.18  0x0                                       No
0b.10.0   0x0                                       Yes
0b.10.12  0x0                                       Yes
0b.10.3   0x0                                       No
0b.10.15  0x0                                       No
0a.00.1   0x0                                       Yes
0a.00.2   0x0                                       Yes
```

- Obtain an IP address for the external key management server.  
See the *Data ONTAP Software Setup Guide for 7-Mode* for more information about the external key management server.
- Take one of the following actions:

If you will use the root volume of the...	Then...
New system after the upgrade	Go to Step 26.
<b>Note:</b> If your system is in a SAN environment, you must use the original root volume after the upgrade.	
Original system after the upgrade, and the root volume is a traditional root volume	Go to Step 26.
Original system after the upgrade, and the root volume is a FlexVol root volume	Go to Step 24.

- Check the size of the FlexVol root volume of the original system by entering one of the following commands:

```
vol size vol_name
```

```
df vol_name
```

### Example

The commands in the following example check the size of a root volume name `vol0`:

```
vol size vol0
```

```
df vol0
```

The FlexVol root volume of the original system must meet the minimum root FlexVol size requirement for the new system before you upgrade the hardware. See the *Hardware Universe* on the NetApp Support Site for the minimum FlexVol sizes for the new system.

- If the FlexVol root volume of the original system is less than the minimum for the new system, increase it by entering the following command:

```
vol size vol_name sizeg
```

In the preceding command, `size` specifies the desired size, and `g` specifies gigabytes.

**Note:** You might need to add disks to the aggregate before you can increase the size of the FlexVol root volume.

### Example

The command in the following example increase the size of a volume named vol0 to 250 gigabytes:

```
vol size vol0 250g
```

Alternatively, you can enter the following command to increase the volume size by a specific amount:

```
vol size vol_name +sizeg
```

In the preceding command, + specifies that the current volume size is to be increased, *size* specifies by how much, and *g* specifies gigabytes.

### Example

The command in the following example increases the size of a volume named vol0 by 20 gigabytes:

```
vol size vol0 +20g
```

For more information about resizing FlexVol volumes, see the *Data ONTAP Storage Management Guide for 7-Mode*.

26. Make sure that all adapters and, if appropriate, the port locations on the original system match the adapter and port locations on the new system.

You might not be able to match the adapter and port locations exactly. If you cannot, note the differences and update the configuration as necessary.

There are two methods for matching the adapter and port locations:

- Running the `setup` command and, on the original system, entering the adapter and port information at the appropriate prompts.  
This is the recommended method. For information about the `setup` command, see the *Data ONTAP Software Setup Guide for 7-Mode*.
- Editing the `/etc/rc` file in the root volume of the original system and then saving the file on the system.  
This method requires that you have an external client and that you are knowledgeable about the file contents.

**Attention:** Take great care in editing the `/etc/rc` file, which contains startup commands for your storage system and network parameters. Errors in the `/etc/rc` file might cause the system to become inaccessible. For information about editing the files in the `/etc` directory, see the *Data ONTAP System Administration Guide for 7-Mode*.

**Note:** You might need to update the `/etc/hosts` file or update host information after entering the `setup` command if the host names contain appended interface IDs.

27. When setup is complete, enter the following command, as directed by the prompt:

```
reboot
```

## Replacing the controller

There are two methods for replacing the controllers, depending on whether you plan to use the root volume of the original system or the root volume of the new system after the upgrade.

### About this task

If you are upgrading a system in an FC or iSCSI environment, or if you are upgrading to a system without internal storage or attached disk shelves, you need to use the original system's root volume on the new system after the upgrade.

### Choices

- [Replacing the controllers using the original system's root volume](#) on page 10

- [Replacing the controllers using the new system's root volume](#) on page 16
- [Changing the default configuration speed for a stretch MetroCluster](#) on page 20
- [Performing post-upgrade tasks](#) on page 25

## Replacing the controllers using the original system's root volume

Replacing the controllers consists of disconnecting the original controllers, removing them from the equipment rack or system cabinet, installing the new controllers, and then moving the disk shelves. You also need to rename the root volume of the new system and—if you are upgrading a system in an HA pair in a stretch MetroCluster—determine if you need to change the default speed of the FC-VI adapter.

### Steps

1. Verify that the original system controllers are in an HA pair by entering the following command on one of the controllers and checking that the output shows that the controller is in HA mode:

```
options cf.mode
```

2. Install the new system and any disk shelves associated with it as described in the *Installation and Setup Instructions* for the system.

The *Installation and Setup Instructions* are included in the controller box and are available on the the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com).

**Note:** Different systems have different numbers of ports. You might need to add a PCI card to the new controllers to match the required configuration for the upgrade.

3. Take one of the following actions, depending on the version of Data ONTAP installed on the original system:

---

#### If the version of Data ONTAP on the original system is... Then...

---

The same as the version on the new system	Go to Step 4.
---	---------------

---

Different from the version on the new system	On both controllers of the original system, install the version of Data ONTAP that you downloaded in <i>Preparing the site and equipment for the upgrade</i> by entering the following command:
--	---

```
software update url -r -f
```

*url* is the location of the Data ONTAP system files.

**Note:** The `-r` option suppresses the automatic reboot. The `-f` option is required to overwrite the `image.tgz` file.

---

4. If the new system has storage, rename its root volume on both of the new system's controllers by using one of the following methods, depending on the type of root volume on the system:

---

#### If the new system has a... Then...

---

Traditional root volume	Rename the root volume to <code>old_vol0</code> by entering the following command:
-------------------------	--

```
vol rename your_root_vol_name old_vol0
```

---

FlexVol root volume	Rename the root volume to <code>old_vol0</code> and the corresponding aggregates to <code>old_aggr0</code> by entering the following commands:
---------------------	--

```
vol rename your_root_vol_name old_vol0
```

```
aggr rename your_aggr_name old_aggr0
```

---

5. If you have not done so already, back up the data on both controllers of the original system, particularly the `/etc/rc` and `/etc/hosts` files, and save any information about site-specific customizations as described in Step 17 of the section *Preparing the site and equipment for the upgrade*.
6. Disable clustering by entering the following command on either controller of the original system and either controller of the new system:  
`cf disable`
7. Shut down the original system and the new system by entering the following command on each system:  
`halt`
8. Properly ground yourself.
9. Take one of the following actions:

If the new system will...	Then...
Be in a SAN environment	<ol style="list-style-type: none"> <li>a. Turn off power to the new system and any disk shelves attached to the system.</li> <li>b. Disconnect any attached disk shelves from the new system.</li> </ol> <p style="text-align: center;"><b>Attention:</b> New disk shelves must not be connected to the new controller before the upgrade is complete or there will be a loss of configuration data.</p>
Not be in a SAN environment	Turn off power to the new system and any disk shelves attached to the system.

10. Turn off the power to the original system and disk shelves attached to the system, and then disconnect the cables.
11. Cable the original system disk shelves to the new system and any disk shelves attached to the new system.  
Follow the instructions in the *Installation and Setup Instructions* for the new system, the appropriate disk shelf guide, and, if you are migrating SAS disk shelves, the *Universal SAS and ACP Cabling Guide*. FC cabling instructions are in the disk shelf guides. Also follow the cabling instructions in the *Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode*.  
**Note:** If you are migrating SAS disk shelves, you should also connect ACP cables. If the new system does not have a dedicated onboard network interface for ACP for each controller, you must dedicate one for each controller at system startup.
12. Turn on the disk shelves and then turn on the new system controllers, observing the startup messages on the console of each controller.
13. When the message `Press Ctrl-C for Boot Menu.` appears, press `Ctrl-C` to display the boot menu on each controller.  
**Note:** Complete Step 14 through Step 32 on both controllers, as appropriate; perform each step on one controller and then on the other, then continue to the next step.
14. Select option **5** at the boot menu and continue with the boot on each controller.  
The Maintenance mode command prompt (`*>`) appears.
15. View disk ownership and determine the new system ID by entering the following command on one of the controllers and examining the output:  
`disk show -v`  
Verify that you can see all the attached storage and make a note of the new system ID. However, if you are in a SAN environment, make sure that no disks from the new system are visible.
16. Verify that you can see all the paths available to each disk by entering the following command and examining its output:  
`storage show disk -p`
17. Assign the disks that you migrated from the original system to the new system by entering the following command:  
`disk reassign -s old_sysid -d new_sysid`

For `old_sysid`, use the System ID number that you captured from the output of the `sysconfig -a` command in Step 6 of the section *Preparing the site and equipment for the upgrade*.

**Note:** If you are reassigning disks from the original system to both controllers in the new system, assign disks from controller A of the original system to controller A of the new system and disks from controller B of the original system to controller B of the new system.

The system displays the following prompts, to which you can respond by entering `y`:

```
Partner node must not be in Takeover mode during disk reassignment from maintenance mode.
Serious problems could result!!
Do not proceed with reassignment if the partner is in takeover mode. Abort reassignment
(y/n)? n

After the node becomes operational, you must perform a takeover and giveback of the HA
partner node to ensure disk reassignment is successful.
Do you want to continue (y/n)? y

Disk ownership will be updated on all disks previously belonging to Filer with sysid
system_id.
Do you want to continue (y/n)? y
```

18. Verify that the disks were reassigned correctly by entering the following command and examining its output:

```
disk show -v
```

19. Verify that the controller and chassis are configured as `ha` by entering the following command and observing the output:

```
ha-config show
```

### Example

The following example shows the output of the `ha-config show` command:

```
*> ha-config show
Chassis HA configuration: ha
Controller HA configuration: ha
```

Systems record in a PROM whether they are in HA pair or stand-alone configuration. The state must be the same on all components within the stand-alone system or HA pair.

If the controller and chassis are not configured as `ha`, use the following commands to correct the configuration:

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

```
ha-config modify chassis ha
```

20. If the new system was used previously, clear the mailbox by entering the following command:

```
mailbox destroy local
```

Mailbox disks store configuration information and, in HA pairs, maintain consistency between nodes and monitor partner activity.

The system displays the following message:

```
Destroying mailboxes forces a node to create new empty mailboxes,
which clears any takeover state, removes all knowledge
of out-of-date plexes of mirrored volumes, and will prevent
management services from going online in 2-node cluster
HA configurations.
Are you sure you want to destroy the local mailboxes? y
```

21. Enter `y`.

22. Exit Maintenance mode by entering the following command at the Maintenance mode prompt:

**halt**

23. Take one of the following actions:

<b>If you are upgrading...</b>	<b>Then...</b>
An HA pair that is not in a Stretch MetroCluster	Skip Step 24 and Step 25 and go to Step 26.
An HA pair that is in a Stretch MetroCluster	Go to Step 24.

24. Use the following table to determine whether you need to reconfigure the default speed (link data rate) of the FC-VI adapters on the new system, depending on how the system will be configured.

The default configuration speed is shown against the `link data rate` output field for an FC-VI adapter in the output of the `sysconfig -v FC-VI adapter_slot` command.

The distance between nodes, the cable type used, and the speed of the FC-VI adapter dictates the default configuration speed of the stretch MetroCluster. By default, FC-VI adapters operate at their highest speed. However, you can achieve a greater supported distance by reducing the speed of the adapter. If you want to change the default speed of the adapter, you need to complete a series of additional steps.

**Note:** You do not need to change the default configuration speed of the FC-VI adapter in a stretch MetroCluster configuration that uses a 16-Gb FC-VI adapter with 16-Gb, long-range SFP connected by a 9/125 μm single-mode fiber (SMF) cable.

**Note:** See *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com) to see what speed of FC-VI adapter that the new controllers support. Not all models of controllers support every speed of adapter.

<b>FC-VI adapter speed in Gbps</b>	<b>Maximum distance in meters</b>						
	<b>4 Gbps SFP</b>		<b>8 Gbps SFP</b>		<b>16 Gbps SFP</b>		
	<b>OM2 cables</b>	<b>OM3, OM3+/OM4 cables</b>	<b>OM2 cables</b>	<b>OM3, OM3+/OM4 cables</b>	<b>OM2 cables</b>	<b>OM3 cables</b>	<b>OM3+/OM4 cables</b>
2	300	500	300	500	N/A	N/A	N/A
4	150	270	150	270	150	270	270
8	N/A	N/A	50	150	50	150	170
16	N/A	N/A	N/A	N/A	35	100	125

25. Take one of the following actions:

<b>If you...</b>	<b>Then...</b>
Need to reset the default speed of the FC-VI adapter	Go to the section <i>Changing the default configuration speed for a stretch MetroCluster</i> .
Do not need to reset the default speed of the FC-VI adapter	Go to Step 28.

26. Boot Data ONTAP by entering the following command at the boot loader prompt:

**boot\_ontap**

27. Take one of the following actions:

<b>If the system you are upgrading to...</b>	<b>Then...</b>
Has storage	Go to Step 29.

If the system you are upgrading to...	Then...
Does not have storage	Skip Step 29 and Step 30 and go to Step 31.

28. Take one of the following actions, depending on the original system root volume:

If the original system had a...	Then...
Traditional root volume	<p>a. Bring the root volume from the original system online and set it to the traditional or FlexVol root volume for the new system by entering the following commands:</p> <pre>vol online vol0 vol options vol0 root</pre> <p>b. Verify that the root volume is online by entering the following command and examining its output:</p> <pre>vol status vol0</pre> <p><b>Note:</b> If the new system does not have any attached disk shelves, vol0 will automatically be online.</p>
FlexVol root volume	<p>a. Bring the aggregate from the original system online and set it to the root volume for the new system by entering the following commands:</p> <pre>aggr online aggr0 vol online vol0 vol options vol0 root</pre> <p>b. Verify that the aggregate is online by entering the following command and examining its output:</p> <pre>aggr status aggr0</pre> <p><b>Note:</b> If the new system does not have any attached disk shelves, aggr0 will automatically be online.</p> <p>c. Verify that the root volume is online by entering the following command and examining its output:</p> <pre>vol status vol0</pre> <p><b>Note:</b> If the new system does not have any attached disk shelves, vol0 will automatically be online.</p>

29. Enter the following command at the system prompt:

```
reboot
```

The reboot ensures that the original root volume is used on the upgraded system after the next boot.

30. Take the following actions to install licenses.

- Wait at least five minutes after the command-line interface is up following the reboot.
- Install licenses as needed by entering the following command on both controllers:

```
license add license_key license_key license_key...
```

The `license_key` is the license key that you obtained in the section *Preparing the site and equipment for the upgrade*. The license key consists of 28 upper-case alphabetic characters. You can add one license key at a time, or you can add multiple license keys at the same time, each separated by a comma or a space.

**Note:** You might be prompted to set up some features and then reboot the system.

31. Verify that the licenses are properly installed by entering the following command and examining its output:

```
license show
```

You might want to compare the output with the output that you capture in Step 13 of the section *Preparing the site and equipment for the upgrade*.

**Note:** The `license show` command without parameters will display both old and new licenses. Use the command `license show -instance` to display a list of the controller serial number for each license.

32. Ensure that the ports are assigned correctly by entering the `setup` command at the system prompt and then changing information as necessary at the appropriate prompts.

33. Download and install the appropriate version of Data ONTAP by entering the following command:

```
software update url -R -f
```

`url` is the location of the Data ONTAP system files.

The `-R` option requests a reboot, which is required at this point to complete the software update. The `-f` option is required to overwrite the `image.tgz` file. For details about installing Data ONTAP, see the *Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode*.

34. (Optional) If `vol0` or `aggr0` were renamed in Step 3, and you plan to remove them, copy the data from the old root volume to the new system's disks.

This step preserves the data from the old root volume if you remove the old root volume in the optional next step. See the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode* for information about copying data.

**Note:** If you copy data from the old root volume, you first must bring the root volume online.

35. (Optional) If `vol0` or `aggr0` were renamed in Step 4, remove them by entering one of the following sets of commands.

If the old root volume on the new system was a...	Then...
Traditional volume	Remove the old root volume by entering the following commands: <pre>vol offline old_vol0 vol destroy old_vol0</pre>
FlexVol volume	Remove the old root volume by entering the following commands: <pre>vol offline old_vol0 vol destroy old_vol0 aggr offline old_aggr0 aggr destroy old_aggr0</pre>

36. Modify the RAID groups and volumes as needed as described in the *Storage Management Guide*.

37. Enable clustering by entering the following command on either controller:

```
cf enable
```

38. Verify that the original system controllers are in an HA pair by entering the following command on one of the controllers and ensuring that the output shows that the controller is in HA mode:

```
cf status
```

39. Configure the SP by using the `system node service-processor network modify` command.

See the *Data ONTAP System Administration Guide for 7-Mode* for information about the SP and the *Data ONTAP Commands: Manual Page Reference for 7-Mode, Volume 1* for detailed information about the `system node service-processor network modify` command.

40. Take one of the following actions:

If the new system is...	Then...
In a SAN environment	<ol style="list-style-type: none"> <li>a. Attach any new disk shelves to the new system, following the installation and setup instructions for the new system and the appropriate disk shelf guide. Also follow the cabling instructions in the <i>High-Availability Configuration Guide</i>.</li> <li>b. Go to the subsection <i>After you finish</i>.</li> </ol>
Not in a SAN environment	Go to the section <i>Performing post-upgrade tasks</i> .

### After you finish

If the system is in a SAN environment, verify that the setup is correct by completing the following substeps:

1. Display the LUNs, WWNNs, and WWPNS by entering the following commands:

```
lun show
```

```
fcp config
```

2. Verify that the LUNs, WWNNs, and WWPNS are unchanged by comparing the output of the same commands that you entered in the section *Preparing the site and equipment for the upgrade*.

If any LUNs, WWNNs, or WWPNS have changed, contact technical support.

3. Verify the initiator/target mode configuration of the FC FCoE cards by completing the following substeps:

- a. Enter the following commands and record their output:

```
fcadmin config
```

```
ucadmin show
```

- b. Verify that the configuration is unchanged by comparing the output of the commands that you entered in Step 3a with the output of the same commands that you entered in the section *Preparing the site and equipment for the upgrade*.

- c. If the mode of any ports have changed, contact technical support.

## Replacing the controllers using the new system's root volume

Replacing the controllers consists of disconnecting the original controllers, removing them from the equipment rack or system cabinet, and then installing the new controllers. You also need to rename the root volume of the original system and—if you are upgrading a system in an HA pair in a stretch MetroCluster configuration—determine if you need to change the default speed of the FC-VI adapter.

### Steps

1. Verify that the original system controllers are in an HA pair by entering the following command on one of the controllers and checking that the output shows that the controller is in HA mode:

```
options cf.mode
```

2. Install the new system and any disk shelves associated with it, as described in the *Installation and Setup Instructions* for that system.

The *Installation and Setup Instructions* are included in the controller box and are available on the the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com).

**Note:** Different systems have different numbers of ports. You might need to add a PCI card to the new controllers to match the required configuration for the upgrade.

3. Take one of the following actions, depending on the version of Data ONTAP that is installed on the original system:

---

<b>If the version of Data ONTAP on the original system is...</b>	<b>Then...</b>
--	----------------

---

The same as the version on the new system	Go to Step 4.
---	---------------

---

Different from the version on the new system	On both controllers of the original system, install the version of Data ONTAP that you downloaded in <i>Preparing for the upgrade</i> by entering the following command:
--	--

```
software update url -r -f
```

*url* is the location of the Data ONTAP system files (.tar, .tgz, or .zip file).

**Note:** The `-r` option suppresses the automatic reboot. The `-f` option is required to overwrite the `image.tgz` file.

---

4. On both controllers of the original system, rename the root volume using one of the following methods, depending on the type of root volume on the system:
- 

<b>If the original system has a...</b>	<b>Then...</b>
--	----------------

---

Traditional root volume	Rename the root volume by entering the following command:
-------------------------	---

```
vol rename your_root_vol_name old_vol0
```

---

FlexVol root volume	Rename the root volume and the corresponding aggregates by entering the following commands:
---------------------	---

```
vol rename your_root_vol_name old_vol0
```

```
aggr rename your_aggr_name old_aggr0
```

---

5. Disable the HA pair by entering the following command on either controller of the original system and on either controller of the new system:

```
cf disable
```

6. Shut down the original system and the new system by entering the following command on the console of each controller:

```
halt
```

7. Properly ground yourself.

8. Turn off the power to the new system and disk shelves.

9. Turn off the power to the original system and its attached disk shelves, and then disconnect the cables.

10. Cable the original system disk shelves to the new system as well as any disk shelves attached to the new system.

Follow the instructions in the *Installation and Setup Instructions* for the new system, the appropriate disk shelf guide, and, if you are migrating SAS disk shelves, the *Universal SAS and ACP Cabling Guide*. FC cabling instructions are in the disk shelf guides. Also follow the cabling instructions in the *Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode*.

**Note:** If you are migrating SAS disk shelves, you should also connect ACP cabling. If the new system does not have a dedicated onboard network interface for ACP for each controller, you must dedicate one for each controller at system startup.

11. Turn on the disk shelves and then turn on the new controllers, observing the startup messages on the console of each of the controller.

12. When the message `Press Ctrl-C for Boot Menu.` appears, press Ctrl-C to display the boot menu on each controller.

**Note:** Complete Step 13 through Step 27, completing each step first on one controller and then on the other controller before moving to the next step.

13. Select option **5** at the boot menu and continue with the boot.

The Maintenance mode command prompt (`*>`) appears.

14. View disk ownership on the new system by entering the following command on one of the controllers:

```
disk show -v
```

Verify that you can see all the attached storage.

15. Verify that you can see all the paths available to each disk by entering the following command and examining its output:

```
storage show disk -p
```

16. Assign the disks that you migrated from the original system to the new system by entering the following command:

```
disk reassign -s old_sysid -d new_sysid
```

For *old\_sysid*, use the System ID number that you captured from the output of the `sysconfig -a` command in Step 6 of the section *Preparing the site and equipment for the upgrade*.

**Note:** If you are reassigning the disks from the original system to both controllers of the new system, assign the disks from controller A of the original system to controller A of the new system and the disks from controller B of the original system to controller B of the new system.

17. Verify that the disks were reassigned correctly by entering the following command and examining its output:

```
disk show -v
```

18. Verify that the original system controllers are in an HA pair by entering the following command on one of the controllers and ensuring that the output shows that the controller is in HA mode:

```
cf status
```

19. If the system was used previously, clear the mailbox by entering the following command:

```
mailbox destroy local
```

Mailbox disks store configuration information and, in HA pairs, maintain consistency between nodes and monitor partner activity.

The system displays the following message:

```
Destroying mailboxes forces a node to create new empty mailboxes,
which clears any takeover state, removes all knowledge
of out-of-date plexes of mirrored volumes, and will prevent
management services from going online in 2-node cluster
HA configurations.
Are you sure you want to destroy the local mailboxes? y
```

20. Enter `y`.

21. Exit Maintenance mode by entering the following command at the Maintenance mode prompt:

```
halt
```

22. Take one of the following actions:

If you are upgrading...	Then...
An HA pair that is not in a Stretch MetroCluster	Skip Step 23 and Step 24 and go to Step 25.
An HA pair that is in a Stretch MetroCluster	Go to Step 23.

23. Use the following table to determine whether you need to reconfigure the default speed (link data rate) of FC-VI adapters on the new system, depending on how the system will be configured.

The default configuration speed is shown against the `link data rate` output field for an FC-VI adapter on running `sysconfig -v FC-VI adapter_slot` command.

The distance between nodes, the cable type used, and the speed of the FC-VI adapter dictates the default configuration speed of the stretch MetroCluster. By default, FC-VI adapters operate at their highest speed. However, you can achieve a greater

supported distance by reducing the speed of the adapter. If you want to change the default speed of the adapter, you need to complete a series of additional steps.

**Note:** You do not need to change the default configuration speed of an FC-VI adapter in a stretch MetroCluster configuration that uses a 16-Gb FC-VI adapter with 16-Gb long-range SFP connected by a 9/125- $\mu$ m single mode fiber (SMF) cable.

FC-VI adapter speed in Gbps	Maximum distance in meters						
	4 Gbps SFP		8 Gbps SFP		16 Gbps SFP		
	OM2 cables	OM3, OM3+/OM4 cables	OM2 cables	OM3, OM3+/OM4 cables	OM2 cables	OM3 cables	OM3+/OM4 cables
2	300	500	300	500	N/A	N/A	N/A
4	150	270	150	270	150	270	270
8	N/A	N/A	50	150	50	150	170
16	N/A	N/A	N/A	N/A	35	100	125

**Note:** See *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com) to see what speed FC-VI adapter that the new controllers support. Not all models of controllers support every speed of adapter.

24. Take one of the following actions:

If you...	Then...
Need to reset the default speed of the FC-VI adapter	Go to the section <i>Changing the default configuration speed for a stretch MetroCluster</i> .
Do not need to reset the default speed of the FC-VI adapter	Go to Step 25.

25. Boot Data ONTAP by entering the following command at the boot environment prompt:

```
boot_ontap
```

26. Check the volume status by entering the following command at the system prompt:

```
vol status
```

27. Optional: Copy data from the old root volume to the new system's disks.

This step preserves the data from the old root volume if you remove the old root volume in the optional next step. See the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode* for information about copying data.

**Note:** If you copy data from the old root volume, you first must bring the root volume online.

28. Optional: Remove the old root volume by entering one of the following sets of commands:

If the old root volume on the original was a...	Then...
Traditional volume	Enter the following commands: <pre>vol offline old_vol0</pre> <pre>vol destroy old_vol0</pre>

If the old root volume on the original was a...	Then...
FlexVol volume	Enter the following commands: <pre> vol offline old_vol0 vol destroy old_vol0 aggr offline old_aggr0 aggr destroy old_aggr0 </pre>

29. Modify RAID groups and volumes as needed as described in the *Data ONTAP Storage Management Guide for 7-Mode*.

30. Take the following actions to install licenses:

- a. Wait at least five minutes after the command-line interface is up following the reboot.
- b. Install licenses as needed by entering the following command on both controllers:

```
license add license_key license_key license_key...
```

The `license_key` is the license key that you obtained in the section *Preparing the site and equipment for the upgrade*. The license key consists of 28 upper-case alphabetic characters. You can add one license key at a time, or you can add multiple license keys at the same time, each separated by a comma or space.

**Note:** You might be prompted to set up some features and then reboot the system.

31. Verify that the licenses are properly installed by entering the following command and examining its output:

```
license show
```

You might want to compare the output with the output that you captured in Step 13 of the section *Preparing the site and equipment for the upgrade*.

**Note:** The `license show` command without parameters will display both old and new licenses. Use the command `license show -instance` to display a list of the controller serial numbers for each license.

32. Ensure that the ports are assigned correctly by entering the `setup` command at the system prompt and then changing information as necessary at the appropriate prompts.

33. Enable the HA pair by entering the following command on either controller:

```
cf enable
```

34. Verify that the original system controllers are in an HA pair by entering the following command on one of the controllers and ensuring that the output shows that the controller is in HA mode:

```
cf status
```

35. Configure the SP by using the `system node service-processor network modify` command.

See the *Data ONTAP System Administration Guide for 7-Mode* for information about the SP and the *Data ONTAP Commands: Manual Page Reference for 7-Mode, Volume 1* for detailed information about the `system node service-processor network modify` command.

## Changing the default configuration speed for a stretch MetroCluster

If you are upgrading to a FAS32xx, FAS62xx, or FAS80xx system in an HA pair configured in a stretch MetroCluster, you might need to change the default configuration speed of the FC-VI adapter.

### Before you begin

- You should complete the following steps only if you have determined that you need to reset the default speed of the FC-VI adapters by using the information in the tables in one of the preceding sections:
  - *Replacing the controller using the original system root volume.*
  - *Replacing the controller using the new system's root volume*

- The FAS32xx, FAS62xx, or FAS80xx system must be at the boot environment prompt; that is, you must have exited Maintenance mode by entering the `halt` command but not rebooted the controller.

### **About this task**

You might need to change the default speed of traffic over the FC-VI adapter for the following upgrades:

- A FAS31xx system to a FAS32xx or FAS62xx system
- A FAS6040 or FAS6080 system to a FAS62xx system
- A FAS31xx, FAS32xx, FAS6040, FAS6080, or FAS62xx system to a FAS80xx system

If you want to increase the distance between nodes, you must decrease the link data rate for the FC-VI adapter to 8 Gb, 4 Gb or 2 Gb, depending on the adapter and SFP type in use.

### **Steps**

1. If you have not done so already, identify the default configuration speed based on the SFP and MMF cable type in use for SMC configuration by referring to the information provided in the table in Step 22 in the section *Replacing the controllers using the original system's root volume* or in Step 21 in the section *Replacing the controllers using the new system's root volume*.
2. Take one of the following actions, depending upon the default configuration speed to be configured as identified in Step 1:

---

**If the default  
FC-VI adapter  
configuration  
speed will be...**

---

**Then...**

8 Gbps

- a. Set the FC-VI speed to 8 Gb by entering the following command at the boot loader prompt:
- b. Verify that the FC-VI speed is set to 8 Gb by entering the following command at the boot loader prompt:

```
setenv ispfcvi-force-8G-only? true
```

```
printenv ispfcvi-force-8G-only?
```

The system console displays output similar to the following example:

```
Variable name          Value
-----
ispfcvi-force-8G-only? true
```

If the FC-VI link data rate for the controller being upgraded to was set to 2 Gb or 4 Gb previously, then perform the following steps.

- c. Ensure that the FC-VI speed is not set to 2 Gb by entering the following command at the boot loader prompt:
- d. Verify that the FC-VI speed is unconfigured for 2 Gb by entering the following commands at the boot loader prompt:

```
unsetenv ispfcvi-force-2G-only?
```

```
printenv ispfcvi-force-2G-only?
```

The system console displays output similar to the following example:

```
Variable name          Value
-----
ispfcvi-force-2G-only? *** Undefined ***
```

- e. Ensure that the FC-VI speed is not set to 4 Gb by entering the following command at the boot loader prompt:
- f. Verify that the FC-VI speed is unconfigured for 4 Gb by entering the following commands at the boot loader prompt:

```
unsetenv ispfcvi-force-4G-only?
```

```
printenv ispfcvi-force-4G-only?
```

The system console displays output similar to the following example:

```
Variable name          Value
-----
ispfcvi-force-4G-only? *** Undefined ***
```

---

---

**If the default  
FC-VI adapter  
configuration  
speed will be...**

---

**Then...**

4 Gbps

- a. Set the FC-VI speed to 4 Gb by entering the following command at the boot loader prompt:

```
setenv ispfcvi-force-4G-only? true
```

- b. Verify that the FC-VI speed is set to 4 Gb by entering the following command at the boot loader prompt:

```
printenv ispfcvi-force-4G-only?
```

The system console displays output similar to the following example:

```
Variable name      Value
-----
ispfcvi-force-4G-only?  true
```

If the FC-VI link data rate for the controller being upgraded to was set to 2 Gb or 8 Gb previously, then perform the following steps:

- c. Ensure that the FC-VI speed is not set to 2 Gb by entering the following command at the boot loader prompt:

```
unsetenv ispfcvi-force-2G-only?
```

- d. Verify that the FC-VI speed is unconfigured for 2 Gb by entering the following commands at the boot loader prompt:

```
printenv ispfcvi-force-2G-only?
```

The system console displays output similar to the following example:

```
Variable name      Value
-----
ispfcvi-force-2G-only?  *** Undefined ***
```

- e. Ensure that the FC-VI speed is not set to 8 Gb by entering the following command at the boot loader prompt:

```
unsetenv ispfcvi-force-8G-only?
```

- f. Verify that the FC-VI speed is unconfigured for 8 Gb by entering the following commands at the boot loader prompt:

```
printenv ispfcvi-force-8G-only?
```

The system console displays output similar to the following example:

```
Variable name      Value
-----
ispfcvi-force-8G-only?  *** Undefined ***
```

---

---

**If the default  
FC-VI adapter  
configuration  
speed will be...**

---

**Then...**

2 Gbps

- a. Set the FC-VI speed to 2 Gb by entering the following command at the boot loader prompt:
- b. Verify that the FC-VI speed is set to 2 Gb by entering the following command at the boot loader prompt:

```
setenv ispfcvi-force-2G-only? true
```

```
printenv ispfcvi-force-2G-only?
```

The system console displays output similar to the following example:

```
Variable name          Value
-----
ispfcvi-force-2G-only? true
```

If the FC-VI link data rate for the controller being upgraded to was set to 4 Gb or 8 Gb previously, then perform the following steps.

- c. Ensure that the FC-VI speed is not set to 4 Gb by entering the following command at the boot loader prompt:
- d. Verify that the FC-VI speed is unconfigured for 4 Gb by entering the following commands at the boot loader prompt:

```
unsetenv ispfcvi-force-4G-only?
```

```
printenv ispfcvi-force-4G-only?
```

The system console displays output similar to the following example:

```
Variable name          Value
-----
ispfcvi-force-4G-only? *** Undefined ***
```

- e. Ensure that the FC-VI speed is not set to 8 Gb by entering the following command at the boot loader prompt:
- f. Verify that the FC-VI speed is unconfigured for 8 Gb by entering the following commands at the boot loader prompt:

```
unsetenv ispfcvi-force-8G-only?
```

```
printenv ispfcvi-force-8G-only?
```

The system console displays output similar to the following example:

```
Variable name          Value
-----
ispfcvi-force-8G-only? *** Undefined ***
```

---

### After you finish

Return to the procedure that you were completing when you were directed to this section:

- *Replacing the controllers using the original root volume.*  
Resume the section at Step 24.
- *Replacing the controller using the new system's root volume.*  
Resume the section at Step 23.

## Performing post-upgrade tasks

After you have completed the upgrade, you should test takeover and giveback and set up Storage Encryption on the new controllers if their disks are encryption-enabled. You also might need to verify and change any converged network adapter (CNA) port configurations. You should also decommission the old hardware.

### Steps

1. [Sending a post-upgrade AutoSupport message](#) on page 25
2. [Testing takeover and giveback](#) on page 25
3. [Setting up Storage Encryption on the new controllers](#) on page 26
4. [Configuring CNA ports](#) on page 27
5. [Decommissioning the old system](#) on page 29

## Sending a post-upgrade AutoSupport message

After you upgrade the controller, you should send an AutoSupport message.

### About this task

Keeping AutoSupport enabled and sending AutoSupport messages when you make changes in your system can help you identify problems and solutions.

### Step

1. Send an AutoSupport message to NetApp by entering the following command:

```
options autosupport.doit completing_disruptive_controller_head_upgrade
```

## Testing takeover and giveback

After you configure all aspects of your HA pair, you need to verify that it is operating as expected in maintaining uninterrupted access to both nodes' storage during takeover and giveback operations. Throughout the takeover process, the local (or takeover) node should continue serving the data normally provided by the partner node. During giveback, control and delivery of the partner's storage should return to the partner node.

### Steps

1. Check the cabling on the HA interconnect cables to make sure that they are secure.
2. Verify that you can create and retrieve files on both nodes for each licensed protocol.
3. Enter the following command from the local node's console:

```
cf takeover
```

See the man page for command details.

The local node takes over the partner node and displays the following output:

```
...
Failover monitor: takeover completed
...
This node has taken over node2.
...

node1(takeover)>
```

**Note:** The local node prompt reflects the takeover state.

4. Check if the controller is in takeover mode by entering the following command:

```
cf status
```

In takeover mode, a controller actively takes over for its partner and displays a message similar to the following:

```
node1 (takeover)> cf status
local_node has taken over partner_node
node1 (takeover)>
```

Additional lines in the output may indicate the reason for the failover or indicate that the partner node is ready for giveback.

5. Use the `sysconfig -r` command to ensure that the local (takeover) node can access its partner's disks.
6. Give back the partner node's data service after it displays the `Waiting for giveback` message by entering the following command:

```
cf giveback
```

The local node releases the partner node, which reboots and resumes normal operation. The following message is displayed on the local node's console when the process is complete:

```
...
Failover monitor: giveback completed
...

NodeA>
```

**Note:** The local node prompt returns to normal.

7. Proceed depending on whether you saw the message that giveback was completed successfully:

<b>If takeover and giveback...</b>	<b>Then...</b>
Is completed successfully	Repeat Step 2 through Step 5 on the partner node.
Fails	Correct the takeover or giveback failure and then repeat this procedure.

## Setting up Storage Encryption on the new controllers

If the new controllers have Storage Encryption enabled, you might need to complete a series of additional steps to ensure uninterrupted Storage Encryption functionality. These steps include collecting network information, obtaining private and public SSL certificates, and running the Storage Encryption setup wizard.

### Before you begin

All the disks on the storage system must be encryption-enabled before you set up Storage Encryption on the new controllers.

### About this task

You can skip this section if the system that you upgraded to does not have Storage Encryption enabled.

If you used Storage Encryption on the original system and migrated the disk shelves to the new system, you can reuse the SSL certificates that are stored on migrated disk drives for Storage Encryption functionality on the upgraded system. However, you should check that the SSL certificates are present on the migrated disk drives. If they are not present you will need to obtain them.

**Note:** Step 1 through Step 3 are only the overall tasks required for configuring Storage Encryption. You need to follow the detailed instructions for each task in the *Data ONTAP Software Setup Guide for 7-Mode*.

### Steps

1. Take one of the following actions:

If the original controller...	Then...
Used Storage Encryption	<p><b>a.</b> Check whether SSL certificates are stored in the <code>/etc/keymgr/cert</code> directory on the disk drives migrated from the original system.</p> <p>You need to have the public and private SSL certificates for the storage system and a public SSL certificate for each external key manager that will be used.</p> <p>See the <i>Data ONTAP System Administration Guide for 7-Mode</i> for information about accessing the <code>/etc</code> directory. Administration information for individual nodes is contained in the 7-Mode guide.</p> <p><b>b.</b> Take one of the following actions:</p> <ul style="list-style-type: none"> <li>• If the SSL certificates are in the <code>/etc/keymgr/cert</code> directory, go to Step 3</li> <li>• If the SSL certificates are not in the <code>/etc/keymgr/cert</code> directory, go to Step 2.</li> </ul>
Did not use Storage Encryption	Go to Step 2.

2. Obtain and install private and public SSL certificates for the storage system and a private SSL certificate for each key management server that you plan to use.
 

Requirements for obtaining the certificates and instructions for installing them are contained in the *Data ONTAP Software Setup Guide for 7-Mode*.
3. Collect the information required to configure Storage Encryption on the new controllers.
 

This includes the network interface name, the network interface IP address, and the IP address for external key management server. The required information is contained in the *Data ONTAP Software Setup Guide for 7-Mode*.
4. Launch and run the Storage Encryption setup wizard, responding to the prompts as appropriate.
5. If you have not done so, repeat Step 2 through Step 4 on the other controller.

#### After you finish

See the *Data ONTAP Storage Management Guide for 7-Mode* for information about managing Storage Encryption on the updated system.

## Configuring CNA ports

If a node has onboard CNA ports or a CNA card, you must check the configuration of the ports and possibly reconfigure them, depending on how you want to use the upgraded system.

#### Before you begin

You must have the correct SFP+ modules for the CNA ports.

#### About this task

CNA ports can be configured into native Fibre Channel (FC) mode or CNA mode. FC mode supports FC initiator and FC target; CNA mode allows concurrent NIC and FCoE traffic over the same 10GbE SFP+ interface and supports FC target.

**Note:** NetApp marketing materials might use the term *UTA2* to refer to CNA adapters and ports. However, the CLI and product documentation use the term *CNA*.

CNA ports might be on an adapter or onboard the controller and have the following configurations:

- CNA cards ordered when the controller is ordered are configured before shipment to have the personality you request.
- CNA cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard CNA ports on new controllers are configured before shipment to have the personality you request.

However, you should check the configuration of the CNA ports on the node and change them, if necessary.

## Steps

1. Enter Maintenance mode by completing the following steps:

---

**If the version of Data ONTAP on the system is... Then...**

---

Data ONTAP 8.2.0 or 8.2.1	<ol style="list-style-type: none"> <li>a. Enter the following command: <b>boot_ontap</b></li> <li>b. Interrupt the boot by pressing Ctrl-C. The system displays the boot menu.</li> <li>c. Select (5) Maintenance mode boot by entering 5, and then enter <b>y</b> when prompted to continue with the boot.</li> </ol>
---------------------------	--

---

Data ONTAP 8.2.2 or later	Access Maintenance mode by entering the following command:  <b>boot_ontap maint</b>
---------------------------	---

---

2. Check how the ports are currently configured by entering one of the following commands on one of the new controllers:

**ucadmin show**

The system displays output similar to the following example:

```
node*> ucadmin show
Adapter      Current Mode   Current Type   Pending Mode   Pending Type   Status
-----
0e           fc       initiator -        -        online
0f           fc       initiator -        -        online
0g           cna      target  -        -        online
0h           cna      target  -        -        online
0e           fc       initiator -        -        online
0f           fc       initiator -        -        online
0g           cna      target  -        -        online
0h           cna      target  -        -        online
*>
```

You can use the **-c** option to display the supported personality mode and FC4 type, as shown in the following example:

```
node*> ucadmin show -c
Adapter  Mode   FC Supported Types   CNA Supported Types
-----
1a       fc     initiator,target    -
1b       fc     initiator,target    -
2a       fc     initiator,target    -
2b       fc     initiator,target    -
```

3. If the current SFP+ module does not match the desired use, replace it with the correct SFP+ module.
4. Examine the output of the `ucadmin show` command and determine whether the CNA ports have the personality you want.
5. Take one of the following actions:

---

**If the CNA ports ...**

**Then...**

Do not have the personality that you want

Go to Step 6.

Have the personality that you want

Skip Step 6 through Step 8 and go to Step 9.

---

6. If the adapter is in initiator mode, and if the CNA port is online, take the CNA port offline by entering the following command:

**storage disable adapter *adapter-name***

Adapters in target mode are automatically offline in Maintenance mode.

7. If the current configuration does not match the desired use, enter the following commands to change the configuration as needed:

```
ucadmin modify -m fc|cna -t initiator|target adapter_name
```

- -m is the personality mode, fc or 10GbE cna.
- -t is the FC4 type, target or initiator.

**Note:** You need to use FC initiator for tape drives, FlexArray Virtualization systems, and Fabric MetroCluster. You also need to use FC initiator for stretch MetroCluster if you are using a FibreBridge6500N bridge. You need to use FC target for SAN clients.

8. Verify the settings by entering the following command and examining its output:

```
ucadmin show
```

9. Enter the following command:

```
halt
```

The system stops at the boot environment prompt.

10. Enter the following command:

```
boot_ontap
```

11. Verify the settings by entering one of the following commands:

```
ucadmin show
```

### Example

The output in the following examples show that the FC4 type of adapter 1b is changing to initiator and that the mode of adapters 2a and 2b is changing to cna.

```
node> ucadmin show
Adapter      Current Current Pending Pending
Mode        Type    Mode    Type    Status
-----
1a          fc      initiator -        -        online
1b          fc      target  -        initiator online
2a          fc      target  cna      -        online
2b          fc      target  cna      -        online
node>
```

12. Online any target ports by entering the following command:

```
fcpc config adapter_name up
```

It is not necessary to online initiator ports. They automatically come online if they are connected to a device.

13. Cable the port.
14. Repeat Step 1 through Step 13 on the other controller.

## Decommissioning the old system

After upgrading, you can decommission the old system through the NetApp Support Site. Decommissioning the system tells NetApp that the system is no longer in operation and removes it from support databases.

### Steps

1. Go to the NetApp Support Site at [mysupport.netapp.com](https://mysupport.netapp.com) and log in.
2. Click the link **My Installed Products**.

3. On the Installed Products page, enter the serial number of the old system in the form and then click **Go!**

A new page displays information about the controller.

4. Make sure that the information about the controller is correct.

If the information about the controller is...	Then...
Correct...	<ol style="list-style-type: none"><li>a. Select <b>Decommission this system</b> in the Product Tool Site drop-down menu.</li><li>b. Go to Step 5.</li></ol>
Not correct...	<ol style="list-style-type: none"><li>a. Click the feedback link to open the form for reporting the problem.</li><li>b. Fill out and submit the form.</li></ol>

5. On the Decommission Form page, fill out the form and click **Submit**.

## How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to [doccomments@netapp.com](mailto:doccomments@netapp.com). To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

## Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Customer Fitness, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.  
NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.  
NetApp, Inc. NetCache is certified RealSystem compatible.