



Controller Hardware Upgrade Express Guide

For AFF and FAS Systems

January 2021 | 215-11429_2021-01_en-us
doccomments@netapp.com

Contents

Deciding whether to use this guide.....3

Considerations for upgrading controller hardware.....5

Choosing methods for upgrading controller hardware.....7

Upgrading controller hardware by moving storage..... 8

 Preparing for the upgrade when moving storage.....9

 Shutting down the original nodes..... 11

 Removing ownership of disks attached to the new nodes..... 12

 Resetting default configuration on the new nodes..... 13

 Installing the new nodes..... 14

 Setting up the new nodes..... 14

 Optional: Moving internal storage or converting the system to a drive shelf..... 16

 Moving internal drives from an original node..... 17

 Converting an original node to a drive shelf..... 18

 Attaching storage shelves and reassigning disk ownership..... 18

 Restoring the root volume configuration..... 19

 Completing the upgrade.....20

 Completing the upgrade in ONTAP 9.8.....21

 Completing the upgrade in ONTAP 9.7 or earlier..... 25

Upgrading controller hardware by moving volumes.....29

 Preparing for the upgrade when moving volumes..... 30

 Installing the new nodes and joining them to the cluster.....31

 Creating an aggregate and moving volumes to the new nodes..... 31

 Moving non-SAN data LIFs and cluster management LIFs to the new nodes.....32

 Deleting SAN LIFs from the original nodes.....33

 Unjoining the original nodes from the cluster..... 34

 Completing the upgrade.....35

Copyright, trademark, and machine translation.....37

 Copyright..... 37

 Trademark..... 37

 Machine translation..... 37

Deciding whether to use this guide

This guide describes how to upgrade controller hardware of an AFF system or a FAS system in a cluster by moving storage or volumes.

You should use this guide if you want to upgrade controller hardware in the following situations:

Note: The hardware upgrade procedures have been simplified in ONTAP 9.8 with the introduction of the automatic port placement feature.

- Your original and new nodes are compatible and supported.
- You are upgrading a pair of nodes running ONTAP 9.0 or later to a new pair of nodes running the same release.
- You are reusing the IP addresses, network masks, and gateways of the original nodes on the new nodes.
- You plan to upgrade controller hardware by moving storage or moving volumes.
- You are prepared to perform a disruptive procedure if you are upgrading by moving storage. Upgrading by moving volumes is nondisruptive.
- You plan to convert a node of a supported model to a disk shelf, and then attach it to the new nodes.

If you are upgrading a two-node MetroCluster configuration, you must use the procedure in the [MetroCluster® Upgrade, Transition, and Expansion Guide](#).


This guide does not describe how to use aggregate relocation to upgrade controller hardware, which is described in the following guides:

- [Using “system controller replace” Commands to Upgrade Controller Hardware Running ONTAP 9.8](#)
- [Using “system controller replace” Commands to Upgrade Controller Hardware Running ONTAP 9.5 to ONTAP 9.7](#)
- [Using Aggregate Relocation to Manually Upgrade Controller Hardware Running ONTAP 9.8 and Later](#)
- [Using Aggregate Relocation to Manually Upgrade Controller Hardware Running ONTAP 9.7 and Earlier](#)

If you are replacing an individual component, see the field-replaceable unit (FRU) flyer for that component.

Note:

In the steps for upgrading controller hardware by moving storage, the original nodes are called node1 and node2, and the new nodes are called node3 and node4. During the described procedure, node1 is replaced by node3, and node2 is replaced by node4.



The terms node1, node2, node3, and node4 are used only to distinguish between the original and new nodes. When following the procedure, you must substitute the real names of your original and new nodes. However, in reality, the names of the nodes do not change: node3 has the name node1, and node4 has the name node2 after the controller hardware is upgraded.

Related concepts

[*Considerations for upgrading controller hardware*](#) on page 5

To plan for the upgrade, you should familiarize yourself with the general upgrade considerations. If necessary, contact technical support for recommendations and guidance specific to the configuration of your cluster.

[*Choosing methods for upgrading controller hardware*](#) on page 7

In general, how you upgrade the controller hardware depends on the platform models of the original nodes. You upgrade either by moving the storage (a disruptive procedure) or by moving the volumes (a nondisruptive procedure).

Related information

[*NetApp Hardware Universe*](#)

Considerations for upgrading controller hardware

To plan for the upgrade, you should familiarize yourself with the general upgrade considerations. If necessary, contact technical support for recommendations and guidance specific to the configuration of your cluster.

Requirements and limitations

- The procedure for upgrading by moving storage is disruptive. The procedure for upgrading by moving volumes is nondisruptive.
- The new nodes must support the same software version as the original nodes.

[ONTAP 9 Upgrade and Revert/Downgrade Guide](#)

[Find the Upgrade and Revert/Downgrade Guide for your version of Data ONTAP 8](#)

- You must observe the maximum cluster size.

When you upgrade by moving volumes, new nodes are joined to the cluster before the original nodes are removed. Ensure the number of controllers in the cluster does not exceed the supported maximum cluster size during the procedure.

[NetApp Hardware Universe](#)

- When combining different controller platform model in a cluster, you must follow storage platform mixing rules.

[NetApp Hardware Universe](#)

- The new nodes must have enough storage to accommodate storage associated with the original nodes.

[ONTAP 9 Disks and Aggregates Power Guide](#)

[Find a Storage Management Guide for your version of Data ONTAP 8](#)

- The root aggregate size and the number of disks supported by the new system must be equal to or greater than those supported by the original system.

For root aggregate size requirements, see the [NetApp Hardware Universe](#)

- The controllers in a HA pair must be two AFF models or two FAS models.
- You can upgrade by moving volumes or physical storage from one FAS2xxx to a later FAS2xxx if the nodes are in the same cluster. For example, you can move volumes or physical storage from a FAS2552 to a FAS2750.
- If you are upgrading a two-node MetroCluster configuration, see: [Choosing a MetroCluster upgrade or refresh method](#)

Systems with internal storage

Note: The following systems have internal storage: FAS2220, FAS2240-2, FAS2240-4, FAS2520, FAS2552, FAS2554, FAS2620, FAS2650, FAS2720, FAS2750, AFF A200, AFF A220, AFF A700s, AFF A800, AFF A250.

- If your system is not listed above, see the [NetApp Hardware Universe](#) to check if it has internal drives.
- If you have a system with internal storage, the recommended upgrade method is upgrading by moving volumes.
- If you have a system with internal storage, you can convert the system to a drive shelf and attach it to a new node in the same cluster. FAS2220 and FAS2520 controllers cannot be converted to a drive shelf.

Converting a FAS2240 or FAS255x system to a drive shelf is an optional task in the workflow for upgrading by moving storage.

- If you have a system with internal storage or a system with volumes or aggregates on internal SATA drives or SSDs, you can upgrade by transferring the internal storage to a drive shelf that is attached to the new node in the same cluster.

Transferring the internal storage is an optional task in the workflow for upgrading by moving storage.

Situations where additional steps might be needed

- If the new system has fewer slots than the original system, or if it has fewer or different types of ports, you might need to add an adapter to the new system.
[NetApp Hardware Universe](#)
- If the original nodes or new nodes use FlexArray Virtualization software.
[NetApp KB Article: What are the specific steps involved in FlexArray for NetApp controller upgrades and replacements](#)
- If your cluster has SAN hosts you might need to take steps to resolve issues with LUN serial number changes.
[NetApp Knowledgebase Answer 1032849: How to resolve issues during storage controller motherboard replacement and head upgrades with iSCSI and FCP](#)
- If your system uses out-of-band ACP, you might need to migrate from out-of-band ACP to in-band ACP.
[NetApp KB Article 1029778: In-Band ACP Setup and Support](#)

Related concepts

[Choosing methods for upgrading controller hardware](#) on page 7

In general, how you upgrade the controller hardware depends on the platform models of the original nodes. You upgrade either by moving the storage (a disruptive procedure) or by moving the volumes (a nondisruptive procedure).

Related tasks

[Upgrading controller hardware by moving storage](#) on page 8

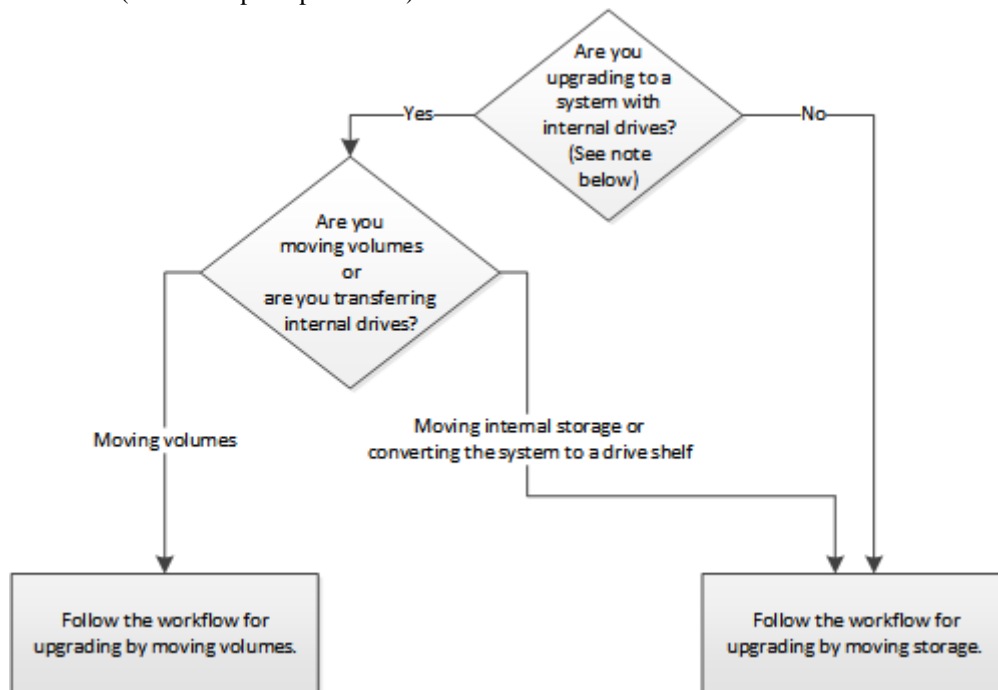
If you are upgrading by moving storage, you prepare the original nodes and set up the new nodes. Some platform models support transferring internal storage to the new nodes. You reassign disks and restore root volume configuration to the new nodes and configure network ports. Upgrading by moving storage is a disruptive procedure.

[Upgrading controller hardware by moving volumes](#) on page 29

If you are upgrading by moving volumes, you prepare the original nodes and join the new nodes to the cluster. You move volumes to the new nodes, configure LIFs, and unjoin the original nodes from the cluster. Upgrading by moving volumes is a nondisruptive procedure.

Choosing methods for upgrading controller hardware

In general, how you upgrade the controller hardware depends on the platform models of the original nodes. You upgrade either by moving the storage (a disruptive procedure) or by moving the volumes (a nondisruptive procedure).



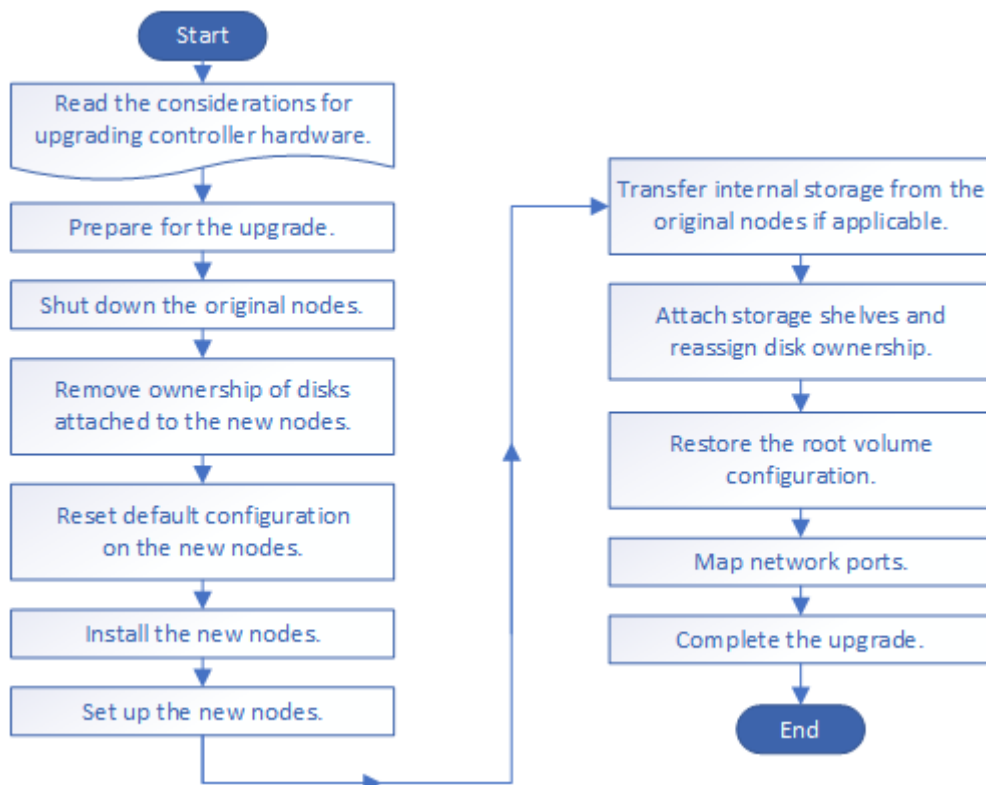
Note: Models with internal drives are: FAS2220, FAS2240-2, FAS2240-4, FAS2520, FAS2552, FAS2554, FAS2620, FAS2650, FAS2720, FAS2750, AFF A200, AFF A220, AFF A700s, AFF A800, AFF A250.

If your system is not listed above, see the [NetApp Hardware Universe](#) to check if it has internal drives.

Upgrading controller hardware by moving storage

If you are upgrading by moving storage, you prepare the original nodes and set up the new nodes. Some platform models support transferring internal storage to the new nodes. You reassign disks and restore root volume configuration to the new nodes and configure network ports. Upgrading by moving storage is a disruptive procedure.

About this task



Steps

1. [Preparing for the upgrade when moving storage](#) on page 9

Before upgrading by moving storage, you must gather license information from the original nodes, plan network configuration, record the system IDs, and prepare the files needed for netboot.

2. [Shutting down the original nodes](#) on page 11

When shutting down and removing the original nodes you must send an AutoSupport message about the upgrade, destroy the mailboxes, power down the nodes, and remove the chassis.

3. [Removing ownership of disks attached to the new nodes](#) on page 12

If the new nodes have internal disks or add-on shelves attached to the system, these can interfere with the controller upgrade. Use the following steps to remove ownership of any new disks that came with node3/node4.

4. [Resetting default configuration on the new nodes](#) on page 13

To ensure that configuration information on the boot media does not interfere with the controller upgrade, you need to reset the configurations of node3 and node4 to the default configuration settings.

5. [Installing the new nodes](#) on page 14

When you upgrade by moving storage, you begin by installing the node3 and node4 and attaching power, console, and network connections to the new nodes.

6. [Setting up the new nodes](#) on page 14

During the process of upgrading by moving storage, you power on node3 and node4, boot the software image, and configure the nodes. Physical port layout between original and new nodes can be different. Mapping of ports between original and replacement nodes should be done to identify proper layout of ports and connections.

7. [Optional: Moving internal storage or converting the system to a drive shelf](#) on page 16

If your original node is one of the supported models, you can move its internal SATA drives or SSDs to a drive shelf that is attached to the new nodes during the process of upgrading by moving storage. You can also convert the system to a drive shelf and attach it to the new nodes.

8. [Attaching storage shelves and reassigning disk ownership](#) on page 18

You must reassign the disks that belonged to node1 and node2 to node3 and node4 respectively.

9. [Restoring the root volume configuration](#) on page 19

You must restore configuration information from the root volume to the boot devices.

10. [Completing the upgrade](#) on page 20

Complete the upgrade in ONTAP 9.8, or ONTAP 9.7 or earlier.

Preparing for the upgrade when moving storage

Before upgrading by moving storage, you must gather license information from the original nodes, plan network configuration, record the system IDs, and prepare the files needed for netboot.

Steps

1. Display and record license information from the original nodes, node1 and node2, by using the `system license show` command.
2. If you use Storage Encryption on the node1/node2 HA pair and the new nodes have encryption-enabled disks, make sure that the original nodes' disks are correctly keyed:
 - a. Display information about self-encrypting disks (SEDs) by using the `storage encryption disk show` command.
 - b. If any disks are associated with a non-manufacture secure ID (non-MSID) key, rekey them to an MSID key by using the `storage encryption disk modify` command.
3. Record port and LIF configuration information on the node1/node2 HA pair:

To display information about...	Enter...
---------------------------------	----------

Shelves, numbers of disks in each shelf, flash storage details, memory, NVRAM, and network cards	<code>system node run -node node_name sysconfig</code>
--	--

Cluster network and node management LIFs	<code>network interface show -role cluster,node-mgmt</code>
--	---

Physical ports	<code>network port show -node node_name -type physical</code>
----------------	---

Failover groups	<code>network interface failover-groups show -vserver vservice_name</code>
-----------------	--

Note: Record the names and ports of failover groups that are not clusterwide.

To display information about...	Enter...
VLAN configuration	<code>network port vlan show -node node_name</code> Note: Record each network port and VLAN ID pairing.
Interface group configuration	<code>network port ifgrp show -node node_name -instance</code> Note: Record the names of the interface groups and the ports assigned to them.
Broadcast domains	<code>network port broadcast-domain show</code>
IPspace information	<code>network ipspace show</code>

4. Obtain information about the default cluster ports, data ports, and node management ports for each new node that you are upgrading to: [NetApp Hardware Universe](#)

5. Download and prepare the files used for performing netboot.

After you install the new nodes, you might need to netboot to ensure the new nodes are running the same version of ONTAP as the original nodes. The term netboot means you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you must put a copy of the ONTAP 9 boot image onto a web server that the system can access.

- a. Access the [NetApp Support Site](#) to download the files used for performing the netboot of the system.
- b. Download the appropriate ONTAP software from the software download section of the NetApp Support Site and store the <ontap_version>_image.tgz file on a web-accessible directory.
- c. Change to the web-accessible directory and verify that the files you need are available.

For...	Then...
FAS2200, FAS2500, FAS3200, FAS6200, FAS/AFF8000 series systems	<p>Extract the contents of the <ontap_version>_image.tgz file to the target directory:</p> <pre>tar -zxvf <ontap_version>_image.tgz</pre> <p>Note: If you are extracting the contents on Windows, use 7-Zip or WinRAR to extract the netboot image.</p> <p>Your directory listing should contain a netboot folder with a kernel file:</p> <pre>netboot/kernel</pre>
All other systems	<p>Your directory listing should contain the following file:</p> <pre><ontap_version>_image.tgz</pre> <p>Note: You do not need to extract the contents of the <ontap_version>_image.tgz file.</p>

You will use information in the directory for [Setting up the new nodes](#) on page 14.

Shutting down the original nodes

When shutting down and removing the original nodes you must send an AutoSupport message about the upgrade, destroy the mailboxes, power down the nodes, and remove the chassis.

Steps

1. Send an AutoSupport message from node1 and node2 to inform technical support of the upgrade:

```
system node autosupport invoke -node node_name -type all -message "MAINT=2h Upgrading
node_name from platform_original to platform_new"
```

2. Disable high availability or storage failover on node1 and node2:

If you have a...	Enter...
Two-node cluster	a. <code>cluster ha modify -configured false</code> b. <code>storage failover modify -node node_name -enabled false</code>
Cluster with more than two nodes	<code>storage failover modify -node node_name -enabled false</code>

3. Halt the node:

```
system node halt -node node_name
```

You can suppress the quorum check during the reboot process by using the `-ignore-quorum-warnings` option.

4. Connect to the serial console, if you are not already connected. The node should be at the LOADER prompt. Use the `boot_ontap maint` command to boot into maintenance mode.

A message might appear asking you to ensure that the partner node is down or takeover is manually disabled on the partner node. You can enter **yes** to continue.

5. Record each original node's system ID, which is obtained through disk ownership information in Maintenance mode:

```
disk show -v
```

You need the system IDs when you assign disks from the original nodes to the new nodes.

```
*> disk show -v
Local System ID: 118049495
DISK   OWNER          POOL   SERIAL NUMBER          HOME
-----
0a.33  node1 (118049495)  Pool0  3KS6BN970000973655KL  node1 (118049495)
0a.32  node1 (118049495)  Pool0  3KS6BCKD000097363ZHK  node1 (118049495)
0a.36  node1 (118049495)  Pool0  3KS6BL9H000097364W74  node1 (118049495)
...
```

6. If you have FC or CNA port configuration, display the configuration in Maintenance mode:

```
ucadmin show
```

You should record the command output for later reference.

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

7. In Maintenance mode, destroy the node1 and node2 mailboxes:

```
mailbox destroy local
```

The console displays a message similar to the following:

```
Destroying mailboxes forces a node to create new empty mailboxes, which  
clears any takeover state, removes all knowledge of out-of-date plexes and  
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?
```

8. Destroy the mailboxes by entering

y

when you see a prompt similar to the following:

```
.....Mailboxes destroyed  
Takeover On Reboot option will be set to ON after the node boots.  
This option is ON by default except on setups that have iSCSI or FCP license.  
Use "storage failover modify -node <nodename> -onreboot false" to turn it OFF.  
*>
```

9. Exit Maintenance mode:

halt

10. Turn off the power to node1 and node2, and then unplug them from the power source.
11. Label and remove all cables from node1 and node2.
12. Remove the chassis containing node1 and node2.

Removing ownership of disks attached to the new nodes

If the new nodes have internal disks or add-on shelves attached to the system, these can interfere with the controller upgrade. Use the following steps to remove ownership of any new disks that came with node3/node4.

About this task

These steps are performed on node3 and node4 one after the other. The node sequence does not matter.

Note:

- The shelves from node1 and node2 are not physically connected to node3 and node4 at this stage.
- You are only required to remove disk ownership for disks and shelves that have come with new controllers.
- You are not required to remove ownership of disks if you are upgrading the hardware by swapping an old controller with a new controller on an internal drive platform while retaining the chassis and disks of the old controller.

For example, if you are upgrading your system from A200 to A220 by only swapping the old A200 controller module with the new A220 controller module while keeping the chassis and disks of the old A200 in place, you would not remove ownership of disks for the new A220 controller module as outlined in this section *Removing ownership of disks attached to the new nodes*.

Contact NetApp technical support if you have questions on removing disk ownership during a controller upgrade.

The following is a list of systems that have internal storage: FAS2220, FAS2240-2, FAS2240-4, FAS2520, FAS2552, FAS2554, FAS2620, FAS2650, FAS2720, FAS2750, AFF A200, AFF A220, AFF A700s, AFF A800, AFF A250.

If your system is not listed above, see the [NetApp Hardware Universe](#) to check if it has internal drives.

Steps

1. At the LOADER prompt of the node, enter the command:

```
boot_ontap menu
```

2. At the boot menu prompt, enter 9a and then press Enter.

The following screen shows the boot menu prompt.

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

```
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
Selection (1-9)? 9a
```

3. Remove disk ownership by entering **y** when you see a prompt similar to the following:

```
##### WARNING #####
```

```
This is a disruptive operation and will result in the
  loss of all filesystem data. Before proceeding further,
  make sure that:
  1) This option (9a) has been executed or will be executed
  on the HA partner node, prior to reinitializing either
  system in the HA-pair.
  2) The HA partner node is currently in a halted state or
  at the LOADER prompt.
```

```
Do you still want to continue (yes/no)? yes
```

The system removes disk ownership and returns to the boot menu.

4. At the boot menu, enter 5 to go to maintenance mode.
5. In maintenance mode, run the `disk show` command.
No disks should be listed.
6. Run the command:

```
disk show -a
```

All listed disks should be unassigned.

7. Exit maintenance mode:

```
halt
```

Resetting default configuration on the new nodes

To ensure that configuration information on the boot media does not interfere with the controller upgrade, you need to reset the configurations of node3 and node4 to the default configuration settings.

About this task

You need to perform the following steps on node3 and node4. You can perform the steps on each node in parallel.

Steps

1. Boot the node to the boot menu:

```
boot_ontap menu
```

2. At the boot menu prompt, type

```
wipeconfig
```

and then press Enter.

The following screen shows the boot menu prompt

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

```
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
Selection (1-9)? wipeconfig
```

3. Enter **yes** when you see a prompt similar to the following:

```
This option deletes critical system configuration, including cluster membership.
Warning: do not run this option on a HA node that has been taken over.
Are you sure you want to continue?: yes
Rebooting to finish wipeconfig request.
```

The system will initiate the wipeconfig procedure and reboot. When the procedure is complete, the system returns to the boot menu.

4. From the boot menu, enter 8 to reboot the node, and press **Ctrl-C** during autoboot to stop the node at the LOADER prompt.

Installing the new nodes

When you upgrade by moving storage, you begin by installing the node3 and node4 and attaching power, console, and network connections to the new nodes.

Steps

1. If needed, install any adapters in node3 and node4, following the instructions in the appropriate adapter installation guide.
2. Install the new nodes, following the *Installation and Setup Instructions* for the platform.
Do not attach disk shelves from the original nodes to the new nodes at this point.
3. Attach power and console connections to the node3/node4 HA pair, following the *Installation and Setup Instructions* for the platform.
4. Attach the network cables.
5. Transfer all remaining cables, other than storage shelf cables, from the node1/node2 HA pair to corresponding ports on node3/node4 respectively.

This includes Fibre Channel and Ethernet cables that are not used to attach storage shelves.

Setting up the new nodes

During the process of upgrading by moving storage, you power on node3 and node4, boot the software image, and configure the nodes. Physical port layout between original and new nodes can

be different. Mapping of ports between original and replacement nodes should be done to identify proper layout of ports and connections.

Before you begin

If the version of ONTAP running on the new nodes is different to the version on the original nodes, you must have downloaded the correct <ontap_version>_image.tgz file from the NetApp Support Site to a web-accessible directory (see Step 5 on page 10). You need the <ontap_version>_image.tgz file to perform a netboot of your system.

Steps

- 1. Turn on the power to node3, and then immediately press Ctrl-C at the console terminal to access the LOADER prompt.
If node3 and node4 are in the same chassis, go to Step 2. If not, go to Step 3.
- 2. If node3 and node4 are in a single-chassis configuration (with controllers in the same chassis):
 - a. Attach a serial console to node4.
 - b. Turn on the power to node4, if it is not already ON, and then interrupt the boot process by pressing Ctrl-C at the console terminal to access the LOADER prompt.
The power should already be ON if both controllers are in the same chassis.
Leave node4 at the LOADER prompt; you return to this procedure and repeat these steps after node3 is installed.

- 3. At the LOADER prompt, enter the following command:

set-defaults

- 4. At the LOADER prompt, configure the netboot connection for a management LIF:

If IP addressing is...	Then...
DHCP	Configure the automatic connection: ifconfig e0M -auto
Static	Configure the manual connection: ifconfig e0M -addr=ip_addr -mask=netmask -gw=gateway

- 5. At the LOADER prompt, perform netboot on the node3:

For...	Then...
FAS2200, FAS2500, FAS3200, FAS6200, FAS/AFF8000 series systems	netboot http://web_server_ip/path_to_web-accessible_directory/netboot/kernel
All other systems	netboot http://web_server_ip/path_to_web-accessible_directory/<ontap_version>_image.tgz

The *path_to_the_web-accessible_directory* is the location of the downloaded <ontap_version>_image.tgz file.

Note: If you are unable to netboot the new controllers, contact technical support.

- 6. From the boot menu, select option (7) **Install new software first** to download and install the new software image to the boot device.

Disregard the following message: "This procedure is not supported for Non-Disruptive Upgrade on an HA pair". It applies to nondisruptive upgrades of software, not to upgrades of controllers.

7. If you are prompted to continue the procedure, enter **y**, and when prompted for the package, enter the URL of the image file:

```
http://web_server_ip/path_to_web-accessible_directory/<ontap_version>_image.tgz
.
```

Enter username/password if applicable, or press Enter to continue.

8. Enter **n** to skip the backup recovery when you see a prompt similar to the following:

```
Do you want to restore the backup configuration now? {y|n}
```

9. Reboot by entering **y** when you see a prompt similar to the following:

```
The node must be rebooted to start using the newly installed software. Do you want to reboot now? {y|n}
```

10. Interrupt the reboot process by pressing Ctrl-C to display the boot menu when the system prompts you to do so.
11. From the boot menu, select **(5) Maintenance mode boot** to access Maintenance mode.
12. If necessary, make changes to the FC or CNA ports on the node, and then reboot the node to Maintenance mode.

[ONTAP 9 SAN Administration Guide](#)

[Find the SAN Administration Guide for your version of Data ONTAP 8](#)

13. You should verify that the command output shows **ha**:

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

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

The `ha-config modify controller ha` command configures **ha** for the controller setting. The `ha-config modify chassis ha` command configures **ha** for the chassis setting.

[ONTAP 9 High-Availability Configuration Guide](#)

[Find the High-Availability Configuration Guide for your version of Data ONTAP 8](#)

14. Exit Maintenance mode:

```
halt
```

The system stops at the LOADER prompt.

Optional: Moving internal storage or converting the system to a drive shelf

If your original node is one of the supported models, you can move its internal SATA drives or SSDs to a drive shelf that is attached to the new nodes during the process of upgrading by moving storage. You can also convert the system to a drive shelf and attach it to the new nodes.

About this task

You can move drives or drive shelves only within the same cluster.

Choices

- [Moving internal drives from an original node](#) on page 17

If your original node is one of the supported models, during the process of upgrading by moving storage you can move the node's internal SATA drives or SSDs to a drive shelf that is attached to the new node in the same cluster. You cannot transfer SAS drives to a drive shelf attached to the new nodes.

- [Converting an original node to a drive shelf](#) on page 18

If your original node is one of the supported models, during the process of upgrading by moving storage you can convert the node to a drive shelf and then attach it to the new nodes in the same cluster.

Moving internal drives from an original node

If your original node is one of the supported models, during the process of upgrading by moving storage you can move the node's internal SATA drives or SSDs to a drive shelf that is attached to the new node in the same cluster. You cannot transfer SAS drives to a drive shelf attached to the new nodes.

Before you begin

- You must have reviewed [Considerations for upgrading controller hardware](#) on page 5 about moving internal drives.
You should contact technical support if you need guidance specific to your configuration.
- The SATA or SSD drive carriers from the original node must be compatible with the new drive shelf.
- A compatible drive shelf must already be attached to the new node.
- The drive shelf must have enough free bays to accommodate the SATA or SSD drive carriers from the original node.

Steps

1. Gently remove the bezel from the front of the system.
2. Press the release button on the left side of the drive carrier.
The cam handle on the carrier partially springs open, and the carrier releases from the midplane.
3. Pull the cam handle to its fully open position to unseat the carrier from the midplane, and then gently slide the carrier out of the drive shelf.



Attention: Always use two hands when removing, installing, or carrying a drive. However, do not place your hands on the drive boards exposed on the underside of the carrier.

4. With the cam handle in the open position, insert the carrier into a slot in the new drive shelf, firmly pushing until the carrier stops.



Attention: Use two hands when inserting the carrier.

5. Close the cam handle so that the carrier is fully seated in the midplane and the handle clicks into place.

You should close the handle slowly so that it aligns correctly with the face of the carrier.

6. Repeat Step 2 on page 17 through Step 5 on page 17 for all of the drives that you are moving to the new system.

Converting an original node to a drive shelf

If your original node is one of the supported models, during the process of upgrading by moving storage you can convert the node to a drive shelf and then attach it to the new nodes in the same cluster.

Before you begin

You must have reviewed [Considerations for upgrading controller hardware](#) on page 5 about converting a node to a drive shelf. You should contact technical support if you need guidance specific to your configuration.

Steps

1. Replace the controller modules in the node you are converting with appropriate IOM modules.
[NetApp Hardware Universe](#)
2. Set the drive shelf ID.
Each drive shelf, including the chassis, requires a unique ID.
3. Reset other drive shelf IDs as needed.
4. Turn off power to any drive shelves connected to the new nodes, and then turn off power to the new nodes.
5. Cable the converted drive shelf to a SAS port on the new system, and, if you are using out-of-band ACP cabling, to the ACP port on the new node.
6. Turn on the power to the converted drive shelf and any other drive shelves attached to the new nodes.
7. Turn on the power to the new nodes, and then interrupt the boot process on each node by pressing Ctrl-C to access the boot environment prompt.

Attaching storage shelves and reassigning disk ownership

You must reassign the disks that belonged to node1 and node2 to node3 and node4 respectively.

About this task

You perform the steps in this section on node3 and node4, completing each step on node3 and then node4 before going on to the next step.

Steps

1. Connect the storage shelf cables from node1/node2 to node3/node4. Verify the power supply and physical connectivity of the shelves.
2. From the node3 LOADER prompt, boot to Maintenance mode:

```
boot_ontap maint
```

3. Display the system ID of node3:

```
disk show -v
```

```
*> disk show -v
Local System ID: 101268854
...
```

Record the system ID of node3 for use in Step 4 below.

4. Reassign node1's spare disks, disks belonging to the root aggregate, and any data aggregates:

```
disk reassign -s node1_sysid -d node3_sysid -p node2_sysID
```

- The parameter *node1_sysid* is the value you recorded in Step 5 on page 11.
- Specify the parameter *-p partner_sysID* only when shared disks are present.

Note: When reassigning node2's spare disks, disks belonging to the root aggregate, and any data aggregates, the command is:

```
disk reassign -s node2_sysid -d node4_sysid -p node3_sysid
```

The system displays a message similar to the following:

```
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
```

5. Enter **y** to continue.

The system displays a message similar to the following:

```
The system displays the following message:
Disk ownership will be updated on all disks previously belonging to Filer with sysid
<sysid>.
Do you want to continue (y/n)? y
```

6. Enter **y** to continue.

7. Ensure that node1's root aggregate is set to root in the options field and that other aggregates are online:

```
aggr status
```

You should see output similar to the following:

```
*> aggr status
      Aggr State      Status      Options
      aggr0 online    raid_dp, aggr    root
                      64-bit
```

8. Exit Maintenance mode:

```
halt
```

Restoring the root volume configuration

You must restore configuration information from the root volume to the boot devices.

About this task

You need to perform these steps on node3 and node4, completing each step on one node and then the other before going on to the next step.

Steps

1. Access the boot menu from the LOADER prompt:

```
boot_ontap menu
```

2. From the boot menu, select (6) Update flash from backup config and enter **y** when prompted to continue.

```
Please choose one of the following:

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

```
This will replace all flash-based configuration with the last backup to  
disks. Are you sure you want to continue?: y
```

The update flash process runs for several minutes, and then the system reboots.

3. When you are prompted to confirm the system ID mismatch, enter **y**.

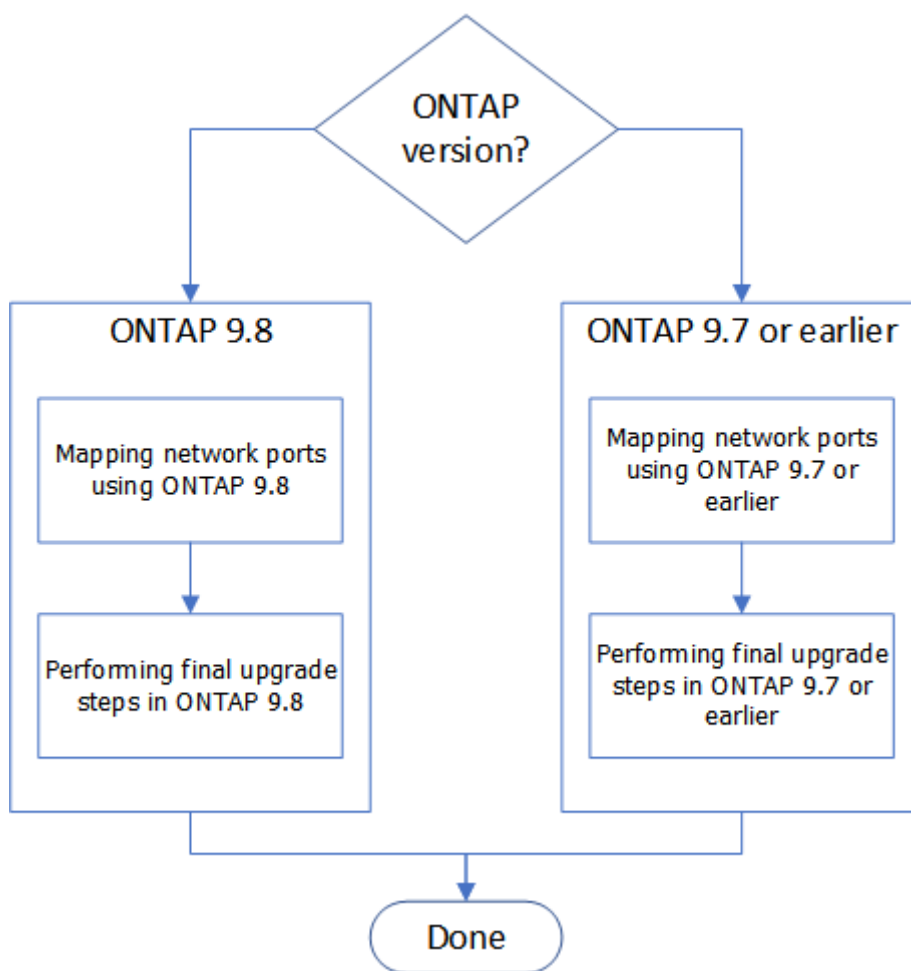
```
WARNING: System id mismatch. This usually occurs when replacing CF or NVRAM cards!  
Override system id? {y|n} [n] y
```

The startup sequence proceeds normally.

Completing the upgrade

Complete the upgrade in ONTAP 9.8, or ONTAP 9.7 or earlier.

You must use the procedure for your version of ONTAP.



- To complete the upgrade in ONTAP 9.8, go to [Completing the upgrade in ONTAP 9.8](#) on page 21.
- To complete the upgrade in ONTAP 9.7 or earlier, go to [Completing the upgrade in ONTAP 9.7 or earlier](#) on page 25.

Completing the upgrade in ONTAP 9.8

Use these steps to complete the upgrade in ONTAP 9.8.

Mapping network ports using ONTAP 9.8

To enable node3 and node4 to communicate with each other in the cluster and with the network after the upgrade, you must ensure the physical ports are correctly configured with the settings for the intended use, such as cluster, data, and so on.

Before you begin

These steps apply to systems running ONTAP 9.8. If you are running ONTAP 9.7 or earlier, you must use the procedure in [Mapping network ports using ONTAP 9.7 or earlier](#) on page 25.

About this task

You must perform these steps on node3 and node4.

Note: The following command examples refer to "node1" because at this stage in the procedure the replacement nodes "node3" and "node4" are actually named "node1" and "node2".

Steps

1. If your system is running ONTAP 9.7 or earlier, **STOP**. You must use the procedure in [Mapping network ports using ONTAP 9.7 or earlier](#) on page 25.
2. Locate the port and LIF configuration information for node1 and node2 that you recorded in Step 3 on page 9.
3. Locate the information for ports, broadcast domains, and IPspaces that you recorded in Step 3 on page 9.
[NetApp Hardware Universe](#)
4. Make the following changes:
 - a. Boot and log in to node3 and node4 if you have not already done so.
 - b. Modify ports that will be part of **Cluster** broadcast domain:

```
network port modify -node node_name -port port_name -mtu 9000 -ipspace Cluster
```

This example adds **Cluster** port e1b on "node1":

```
network port modify -node node1 -port e1b -ipspace Cluster -mtu 9000
```
 - c. Migrate the cluster LIFs to the new ports, once for each LIF:

```
network interface migrate -vserver vserver_name -lif lif_name -  
source-node node1 -destination-node node1 -destination-port  
port_name
```

When all cluster LIFs are migrated and cluster communication is established, the cluster should come into quorum.
 - d. Modify the home port of the **Cluster** LIFs:

```
network interface modify -vserver Cluster -lif lif_name -home-port  
port_name
```
 - e. Remove the old ports from the **Cluster** broadcast domain:

```
network port broadcast-domain remove-ports -ipspace Cluster -broadcast-domain  
Cluster -ports node1:port
```
 - f. Display the **health** state of node3 and node4:

```
cluster show -node node1 -fields health
```
5. Modify the broadcast domain memberships of physical ports hosting data LIFs. You can do this manually, as shown in Step 7 on page 26 of [Mapping network ports using ONTAP 9.7 or earlier](#). The recommended approach is to use the enhanced network reachability scan and

repair procedure introduced in ONTAP 9.8, as shown in substeps (a) to (g) of Step 5 on page 21.

- a. List the reachability status of all ports:

```
network port reachability show
```

- b. Repair the reachability of the physical ports, followed by VLAN ports, by running the following command on each port, one port at a time:

```
reachability repair -node node_name -port port_name
```

A warning like the following is expected. Review and enter y or n as appropriate:

```
Warning: Repairing port "node_name:port" may cause it to move into a
different broadcast domain, which can cause LIFs to be re-homed away
from the port. Are you sure you want to continue? {y|n}:
```

- c. To allow ONTAP to complete the repair, wait for about a minute after running the reachability repair command on the last port.

- d. List all broadcast domains on the cluster:

```
broadcast-domain show
```

- e. As the reachability repair is performed, ONTAP attempts to place the ports in the correct broadcast domains. However, if a port's reachability cannot be determined and does not correspond to any of the existing broadcast domains, ONTAP will create new broadcast domains for these ports. As required, you can delete the newly created broadcast domains if all their member ports will become member ports of the interface groups. Delete broadcast domains:

```
broadcast-domain delete -broadcast-domain broadcast_domain
```

- f. Review the interface group configuration, and as required, add or delete member ports. Add member ports to interface group ports:

```
ifgrp add-port -node node_name -ifgrp ifgrp_port -port port_name
```

Remove member ports from interface group ports:

```
ifgrp remove-port -node node_name -ifgrp ifgrp_port -port port_name
```

- g. Delete and re-create VLAN ports as needed. Delete VLAN ports:

```
vlan delete -node node_name -vlan-name vlan_port
```

Create VLAN ports:

```
vlan create -node node_name -vlan-name vlan_port
```

Note: Depending on the complexity of the networking configuration of the system being upgraded, you might be required to repeat substeps (a) to (g) of Step 5 on page 21 until all ports are placed correctly where needed.

6. If there are no VLANs configured on the system, go to Step 7 on page 23. If there are VLANs configured, restore displaced VLANs that were previously configured on ports that no longer exist or were configured on ports that were moved to another broadcast domain.

- a. Display the displaced VLANs:

```
displaced-vlans show
```

- b. Restore the displaced VLANs to the desired destination port:

```
displaced-vlans restore -node node_name -port port_name -destination-port
destination_port
```

- c. Verify that all displaced VLANs have been restored:

```
displaced-vlans show
```

- d. VLANs are automatically placed into the appropriate broadcast domains about a minute after they are created. Verify that the restored VLANs have been placed into the appropriate broadcast domains:

network port reachability show

7. Starting from ONTAP 9.8, ONTAP will automatically modify the home ports of LIFs if the ports are moved between broadcast domains during the network port reachability repair procedure. If a LIF's home port was moved to another node, or is unassigned, that LIF will be presented as a displaced LIF. Restore the home ports of displaced LIFs whose home ports either no longer exist or were relocated to another node.

- a. Display the LIFs whose home ports might have moved to another node or no longer exist:

displaced-interface show

- b. Restore the home port of each LIF:

displaced-interface restore -vserver vservice_name -lif-name lif_name

- c. Verify that all LIF home ports have been restored:

displaced-interface show

When all ports are correctly configured and added to the correct broadcast domains, the `network port reachability show` command should report the reachability status as `ok` for all connected ports, and the status as `no-reachability` for ports with no physical connectivity. If any ports are reporting a status other than these two, repair the reachability as outlined in Step 5 on page 21.

8. Verify that all LIFs are administratively up on ports belonging to the correct broadcast domains.

- a. Check for any LIFs that are administratively down:

network interface show -vserver vservice_name -status-admin down

- b. Check for any LIFs that are operationally down:

network interface show -vserver vservice_name -status-oper down

- c. Modify any LIFs that need to be modified to have a different home port:

network interface modify -vserver vservice_name -lif lif -home-port home_port

Note: For iSCSI LIFs, modification of the home port requires the LIF to be administratively down.

- d. Revert LIFs that are not home to their respective home ports:

network interface revert *

After you finish

You have completed mapping the physical ports. To complete the upgrade, go to [Performing final upgrade steps in ONTAP 9.8](#) on page 23.

Performing final upgrade steps in ONTAP 9.8

To complete the procedure of upgrading by moving storage, you must delete any unused ports and LIFs from the new nodes, re-enable storage failover or high availability, configure the Service Processor (SP), install new licenses, and set up AutoSupport. You might also need to set up Storage or Volume Encryption and configure the FC or CNA ports.

Before you begin

These steps apply to systems running ONTAP 9.8. If you are running ONTAP 9.7 or earlier, you must use the procedure in [Performing final upgrade steps in ONTAP 9.7 or earlier](#) on page 26.

Steps

1. If your system is running ONTAP 9.7 or earlier, **STOP**. You must use the procedure in [Performing final upgrade steps in ONTAP 9.7 or earlier](#) on page 26.
2. From the storage system prompt, display information about LIFs:

network interface show

3. If you are in a SAN environment, delete unused LIFs from the port set so that you can remove them:

a. Display the port set list:

lun portset show

b. Remove any unused LIFs from the port set:

lun portset remove

4. Remove each unused LIF from the new nodes:

network interface delete

5. Re-enable storage failover or high availability on the new node pair as needed:

If you have a...	Then...
Two-node cluster	Re-enable high availability: cluster ha modify -configured true
A cluster with more than two nodes	Re-enable storage failover: storage failover modify -node node_name -enabled true

6. Configure the SP on the new nodes as needed:

system service-processor network modify

7. Install new licenses on the new nodes as needed:

system license add

[NetApp Knowledgebase Answer 1002749: Data ONTAP 8.2 and 8.3 Licensing Overview and References](#)

8. Set up AutoSupport on the new nodes:

system node autosupport modify

9. From each new node, send a post-upgrade AutoSupport message to technical support:

system node autosupport invoke -node node_name -type all -message "MAINT=END node_name successfully upgraded from platform_old to platform_new"

10. Restore Storage or Volume Encryption functionality by using the appropriate procedure in the *NetApp Encryption Power Guide*.

[ONTAP 9 NetApp Encryption Power Guide](#)

Use one of the following procedures, depending on whether you are using onboard or external key management:

- "Restoring onboard key management encryption keys"
- "Restoring external key management encryption keys"

11. If the new nodes have FC ports (onboard or on FC adapters), onboard CNA ports, or a CNA card, configure the FC or CNA ports by entering the following command at the storage system prompt:

system node hardware unified-connect modify -node node-name -adapter adapter-name -mode {fc|cna} -type {target|initiator}

[ONTAP 9 SAN Administration Guide](#)

[Find the SAN Administration Guide for your version of Data ONTAP 8](#)

You can modify the CNA configuration only when the CNA adapters are offline.

12. Set up a switchless cluster on the new nodes if necessary.

[Migrating to a two-node switched cluster with Cisco cluster switches](#)

[Migrating to a two-node switched cluster with NetApp CN1610 cluster switches](#)

13. As needed, decommission the original systems through the NetApp Support Site to inform NetApp that the systems are no longer in operation and can be removed from support databases:
 - a. Log in to the [NetApp Support](#) site.
 - b. Click the link **My Installed Systems**.
 - c. On the **Installed Systems** page, enter the serial number of the old system in the form and then click **Go!**
 - d. On the **Decommission Form** page, fill out the form and click **Submit**.

After you finish

You have completed the upgrade procedure.

Completing the upgrade in ONTAP 9.7 or earlier

Use these steps to complete the upgrade in ONTAP 9.7 or earlier.

Mapping network ports using ONTAP 9.7 or earlier

To enable node3 and node4 to communicate with each other in the cluster and with the network after the upgrade, you need to ensure the physical ports are correctly configured with the settings for the intended use, such as cluster, data, and so on.

Before you begin

These steps apply to systems running ONTAP 9.7 or earlier. If you are running ONTAP 9.8, you must use the procedure in [Mapping network ports using ONTAP 9.8](#) on page 21.

About this task

You need to perform these steps on node3 and node4.

Note: The following command examples refer to "node1" because at this stage in the procedure the replacement nodes "node3" and "node4" are actually named "node1" and "node2".

Steps

1. If your system is running ONTAP 9.8, **STOP**. You must use the procedure in [Mapping network ports using ONTAP 9.8](#) on page 21.
2. Locate the port and LIF configuration information for node1 and node2 that you recorded in Step 3 on page 9.
3. Locate the information for ports, broadcast domains, and IPspaces that you recorded in Step 3 on page 9.
[NetApp Hardware Universe](#)
4. Make the following changes:
 - a. Boot node3 and node4 to the cluster prompt if you have not already done so.
 - b. Add the correct ports to the **Cluster** broadcast domain:

```
network port modify -node node_name -port port_name -mtu 9000 -ipspace Cluster
```

This example adds **Cluster** port **e1b** on "node1":

```
network port modify -node node1 -port e1b -ipspace Cluster -mtu 9000
```
 - c. Migrate the LIFs to the new ports, once for each LIF:

```
network interface migrate -vserver vservers_name -lif lif_name -source-node node1 -destination-node node1 -destination-port port_name
```

SAN data LIFs can be migrated only when they are offline.

- d. Modify the home port of the **Cluster** LIFs:

```
network interface modify -vserver Cluster -lif lif_name -home-port port_name
```
- e. Remove the old ports from the **Cluster** broadcast domain:

```
network port broadcast-domain remove-ports -ipSpace Cluster -broadcast-domain Cluster -ports node1:port
```
- f. Display the **health** state of node3 and node4:

```
cluster show -node node1 -fields health
```
5. Modify the VLAN and `ifgrp config` to match the new controller physical port layout.
6. Delete the node1 and node2 ports that no longer exist on node3 and node4 (advanced privilege level):

```
network port delete -node node1 -port port_name
```
7. Adjust the node-management broadcast domain and migrate the node-management and cluster-management LIFs if necessary:
 - a. Display the home port of a LIF:

```
network interface show -fields home-node,home-port
```
 - b. Display the broadcast domain containing the port:

```
network port broadcast-domain show -ports node_name:port_name
```
 - c. Add or remove ports from broadcast domains as necessary:

```
network port broadcast-domain add-ports  
network port broadcast-domain remove-ports
```
 - d. Modify a LIF's home port if necessary:

```
network interface modify -vserver vserver_name -lif lif_name -home-port port_name
```
8. Adjust the intercluster broadcast domains and migrate the intercluster LIFs, if necessary, using the commands in Step 7 on page 26.
9. Adjust any other broadcast domains and migrate the data LIFs, if necessary, using the commands in Step 7 on page 26.
10. Adjust all the LIF failover groups:

```
network interface modify -failover-group failover_group -failover-policy failover_policy
```

The following command sets the failover policy to broadcast-domain-wide and uses the ports in failover group "fg1" as failover targets for LIF "data1" on "node1":

```
network interface modify -vserver node1 -lif data1 -failover-policy broadcast-domain-wide -failover-group fg1
```
11. Display node3 and node4 's network port attributes:

```
network port show -node node1
```

After you finish

You have completed mapping the physical ports. To complete the upgrade, go to [Performing final upgrade steps in ONTAP 9.7 or earlier](#) on page 26.

Performing final upgrade steps in ONTAP 9.7 or earlier

To complete the procedure of upgrading by moving storage, you need to delete any unused ports and LIFs from the new nodes, re-enable storage failover or high availability, configure the Service Processor (SP), install new licenses, and set up AutoSupport. You might also need to set up Storage or Volume Encryption and configure the FC or CNA ports.

Before you begin

These steps apply to systems running ONTAP 9.7 or earlier. If you are running ONTAP 9.8, you must use the procedure in [Performing final upgrade steps in ONTAP 9.8](#) on page 23.

Steps

1. If your system is running ONTAP 9.8, **STOP**. You must use the procedure in [Performing final upgrade steps in ONTAP 9.8](#) on page 23.
2. From the storage system prompt, display information about LIFs:
`network interface show`
3. Delete any unused ports from the new nodes (advanced privilege level):
`network port delete`
4. If you are in a SAN environment, delete unused LIFs from the port set so that you can remove them:
 - a. Display the port set list:
`lun portset show`
 - b. Remove any unused LIFs from the port set:
`lun portset remove`
5. Remove each unused LIF from the new nodes:
`network interface delete`
6. Re-enable storage failover or high availability on the new node pair as needed:

If you have a...	Then...
Two-node cluster	Re-enable high availability: <code>cluster ha modify -configured true</code>
A cluster with more than two nodes	Re-enable storage failover: <code>storage failover modify -node node_name -enabled true</code>

7. Configure the SP on the new nodes as needed:
`system service-processor network modify`
8. Install new licenses on the new nodes as needed:
`system license add`
[NetApp Knowledgebase Answer 1002749: Data ONTAP 8.2 and 8.3 Licensing Overview and References](#)
9. Set up AutoSupport on the new nodes:
`system node autosupport modify`
10. From each new node, send a post-upgrade AutoSupport message to technical support:
`system node autosupport invoke -node node_name -type all -message "MAINT=END node_name successfully upgraded from platform_old to platform_new"`
11. Restore Storage or Volume Encryption functionality by using the appropriate procedure in the *NetApp Encryption Power Guide*.

[ONTAP 9 NetApp Encryption Power Guide](#)

Use one of the following procedures, depending on whether you are using onboard or external key management:

- "Restoring onboard key management encryption keys"

- "Restoring external key management encryption keys"
12. If the new nodes have FC ports (onboard or on FC adapters), onboard CNA ports, or a CNA card, configure the FC or CNA ports by entering the following command at the storage system prompt:


```
system node hardware unified-connect modify -node node-name -adapter adapter-name -mode {fc|cna} -type {target|initiator}
```

[ONTAP 9 SAN Administration Guide](#)
[Find the SAN Administration Guide for your version of Data ONTAP 8](#)

You can modify the CNA configuration only when the CNA adapters are offline.
 13. Set up a switchless cluster on the new nodes if necessary.

[Migrating to a two-node switched cluster with Cisco cluster switches](#)
[Migrating to a two-node switched cluster with NetApp CN1610 cluster switches](#)
 14. As needed, decommission the original systems through the NetApp Support Site to inform NetApp that the systems are no longer in operation and can be removed from support databases:
 - a. Log in to the [NetApp Support](#) site.
 - b. Click the link **My Installed Systems**.
 - c. On the **Installed Systems** page, enter the serial number of the old system in the form and then click **Go!**
 - d. On the **Decommission Form** page, fill out the form and click **Submit**.

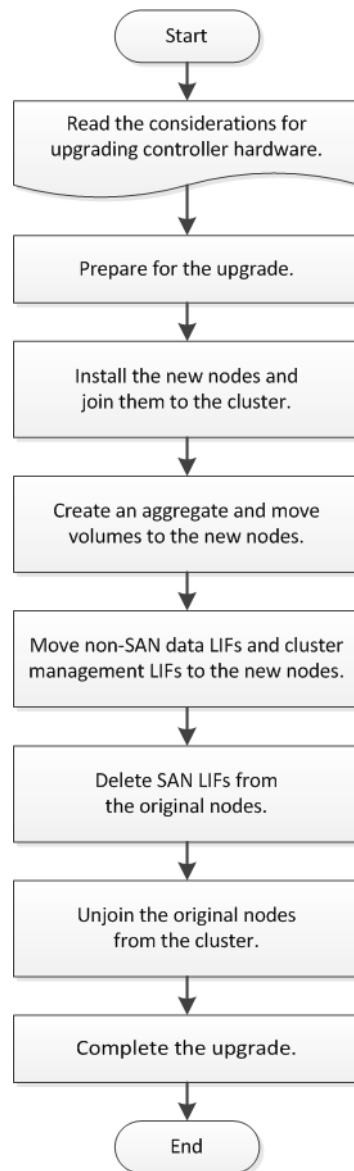
After you finish

You have completed the upgrade procedure.

Upgrading controller hardware by moving volumes

If you are upgrading by moving volumes, you prepare the original nodes and join the new nodes to the cluster. You move volumes to the new nodes, configure LIFs, and unjoin the original nodes from the cluster. Upgrading by moving volumes is a nondisruptive procedure.

About this task



Steps

1. [Preparing for the upgrade when moving volumes](#) on page 30
You need to perform a few preparation steps before upgrading controller hardware by moving volumes.
2. [Installing the new nodes and joining them to the cluster](#) on page 31

You must install the new nodes and join them to the cluster so that you can move volumes from the original nodes.

3. [Creating an aggregate and moving volumes to the new nodes](#) on page 31

You create at least an aggregate on each of the new nodes to store the volumes you want to move from the original nodes. You must identify an aggregate for each volume and move each volume individually.

4. [Moving non-SAN data LIFs and cluster management LIFs to the new nodes](#) on page 32

After you have moved the volumes from the original nodes, you need to migrate the non-SAN data LIFs and cluster-management LIFs from the original nodes to the new nodes.

5. [Deleting SAN LIFs from the original nodes](#) on page 33

If the cluster is in a SAN environment, you must delete any SAN LIFs from the original nodes before you can unjoin the original nodes from the cluster.

6. [Unjoining the original nodes from the cluster](#) on page 34

After the volumes have been moved to the new nodes, you unjoin the original nodes from the cluster. When you unjoin a node, the node's configuration is erased and all disks are initialized.

7. [Completing the upgrade](#) on page 35

To complete the procedure of upgrading by moving volumes, you need to configure the Service Processor (SP), install new licenses, and set up AutoSupport. You might also need to set up Storage or Volume Encryption and configure the FC or NCA ports.

Preparing for the upgrade when moving volumes

You need to perform a few preparation steps before upgrading controller hardware by moving volumes.

Steps

1. Display the volumes on the original nodes by using the `volume show` command.
You use the command output to prepare the list of volumes to move to the new nodes.
2. Display and record license information from the original nodes by using the `system license show` command.
3. If you use Storage Encryption on the original nodes and the new nodes have encryption-enabled disks, make sure that the original nodes' disks are correctly keyed:
 - a. Display information about self-encrypting disks (SEDs) by using the `storage encryption disk show` command.
 - b. If any disks are associated with a non-manufacture secure ID (non-MSID) key, rekey them to an MSID key by using the `storage encryption disk modify` command.
4. If the cluster is currently in a two-node switchless configuration, migrate the cluster to a two-node switched cluster using the type of switch you prefer.
[Migrating to a two-node switched cluster with Cisco cluster switches](#)
[Migrating to a two-node switched cluster with NetApp CN1610 cluster switches](#)
5. Send an AutoSupport message from each original node to inform technical support of the upgrade:

```
system node autosupport invoke -node node_name -type all -message "Upgrading node_name  
from platform_original to platform_new"
```

Installing the new nodes and joining them to the cluster

You must install the new nodes and join them to the cluster so that you can move volumes from the original nodes.

About this task

When you upgrade controller hardware by moving volumes, both the original nodes and the new nodes must be in the same cluster.

Step

Install the new nodes and join them to the cluster:

If the cluster is running...	Follow instructions in...
ONTAP 9.0 or later	ONTAP 9 Cluster Expansion Express Guide
Releases before ONTAP 9.0	Find the Cluster Expansion Express Guide for your version of Data ONTAP 8

Creating an aggregate and moving volumes to the new nodes

You create at least an aggregate on each of the new nodes to store the volumes you want to move from the original nodes. You must identify an aggregate for each volume and move each volume individually.

Before you begin

Data protection mirror relationships must have been initialized before you can move a volume.

[Find a Data Protection Guide for your version of Data ONTAP 8](#)

Steps

1. Create at least one aggregate on each new node:

```
storage aggregate create -aggregate aggr_name -node new_node_name -diskcount integer
```

2. Add the new aggregate to the same SVM as the aggregate on the original node from which you want to move the volumes by using the `vserver add-aggregates` command.

Both the new aggregate and the old aggregate from which the volume will be moved must be in the same SVM.

3. Verify that the new aggregate is now assigned to the same SVM as the aggregate on the original node by using the `vserver show -vserver svm_name` command.
4. Display information for the volumes that you want to move from the original nodes to the new nodes:

```
volume show -vserver svm_name -node original_node_name
```

You should retain the command output for later reference.

The following example displays volumes on the "vs1" SVM and the "node0" node:

```
cluster::> volume show -vserver vs1 -node node0
Vserver  Volume      Aggregate  State    Type    Size    Available  Used%
-----
vs1      clone       aggr1     online   RW      40MB    37.87MB   5%
vs1      vol1        aggr1     online   RW      40MB    37.87MB   5%
vs1      vs1root     aggr1     online   RW      20MB    18.88MB   5%
3 entries were displayed.
```

5. Determine an aggregate to which you can move a given volume:

```
volume move target-aggr show -vserver svm_name -volume vol_name
```

The following example shows that the "user_max" volume on the "vs2" SVM can be moved to any of the listed aggregates:

```
cluster::> volume move target-aggr show -vserver vs2 -volume user_max
Aggregate Name    Available Size  Storage Type
-----
aggr2             467.9GB        FCAL
node12a_aggr3     10.34GB        FCAL
node12a_aggr2     10.36GB        FCAL
node12a_aggr1     10.36GB        FCAL
node12a_aggr4     10.36GB        FCAL
5 entries were displayed
```

6. Run a validation check on each volume that you want to move to verify that it can be moved to the specified aggregate:

```
volume move start -vserver svm_name -volume volume_name -destination-aggregate
destination_aggregate_name -perform-validation-only true
```

7. Move the volumes one at a time (advanced privilege level):

```
volume move start -vserver svm_name -volume vol_name -destination-aggregate
destination_aggr_name -cutover-window integer
```

You cannot move the node root volume (vol0). Other volumes, including SVM root volumes, can be moved.

8. Display the outcome of the `volume move` operation to verify that the volumes were moved successfully:

```
volume move show -vserver svm_name -volume vol_name
```

9. If the `volume move` operation does not complete the final phase after multiple attempts, force the move to finish:

```
volume move trigger-cutover -vserver svm_name -volume vol_name -force true
```

Forcing the `volume move` operation to finish can disrupt client access to the volume that you are moving.

10. Verify that the volumes were moved successfully to the specified SVM and are in the correct aggregate:

```
volume show -vserver svm_name
```

Moving non-SAN data LIFs and cluster management LIFs to the new nodes

After you have moved the volumes from the original nodes, you need to migrate the non-SAN data LIFs and cluster-management LIFs from the original nodes to the new nodes.

About this task

You cannot migrate a LIF that is used for copy-offload operations with VMware vStorage APIs for Array Integration (VAAI).

Steps

1. From the node where the cluster LIF is hosted, change the home ports for the non-SAN data LIFs from the original nodes to the new nodes:

```
network interface modify -vserver vs_server_name -lif lif_name -home-node new_node_name -
home-port {netport|ifgrp}
```

2. Take one of the following actions:

If you want to migrate...	Then enter...
A specific LIF	<code>network interface migrate -vserver vs0 -lif lif_name -source-node source_node_name -destination-node dest_node_name -destination-port dest_port_name</code>
All the non-SAN data LIFs and cluster-management LIFs	<code>network interface migrate-all -node node_name</code>

The following command migrates a LIF named `data1if1` on the SVM `vs0` to the port `e0d` on node0b:

```
cluster::> network interface migrate -vserver vs0 -lif data1if1 -destination-node node0b -destination-port e0d
```

The following command migrates all the data and cluster-management LIFs from the current (local) node:

```
cluster::> network interface migrate-all -node local
```

3. Check whether the home node of the cluster-management LIF is on one of the original nodes:

```
network interface show -lif cluster_mgmt -fields home-node
```

4. If the home node of the cluster management LIF is on one of the original nodes, complete the following steps:

- a. Switch the home node of the cluster-management LIF to one of the new nodes:

```
network interface modify -vserver cluster_name -lif cluster_mgmt -home-node new_node_name -home-port {netport|ifgrp}
```

- b. Migrate the cluster-management LIF to one of the new nodes:

```
network interface migrate -vserver vs0 -lif cluster_mgmt -destination-node new_node_name -destination-port {netport|ifgrp}
```

Deleting SAN LIFs from the original nodes

If the cluster is in a SAN environment, you must delete any SAN LIFs from the original nodes before you can unjoin the original nodes from the cluster.

Steps

1. If you have iSCSI initiators, complete the following steps:

- a. Display a list of active initiators currently connected to an SVM on the original nodes, once for each of the old LIFs:

```
iscsi connection show -vserver Vserver_name -lif old_lif
```

The following example shows the output of the command with an active initiator connected to SVM `vs1`:

```
cluster::> iscsi connection show -vserver vs1 -lif data2
```

Vserver	Tpgroup Name	Conn TSIH	Local ID	Local Address	Remote Address	TCP Recv Size
vs1	data	9	1	10.229.226.166	10.229.136.188	131400

- b. If any initiators are still logged in to an original node, log out of the sessions from your host computer.
2. Display the port set list to determine if any iSCSI or FC LIFs on the original nodes belong to a port set:

```
lun portset show
```

The following example shows output of the `lun portset show` command:

```
cluster:> lun portset show
Virtual
Server      Portset      Protocol Port Names      Igroups
-----
js11        ps0          mixed    LIF1,           igroup1
            ps1          iscsi    LIF2
            ps2          fcp      LIF3
            ps2          fcp      LIF4           igroup2
            ps2          fcp      LIF4           -
3 entries were displayed.
```

3. If any iSCSIs or FC LIFs on an original node are members of a port set, remove them from the port set:

```
lun portset remove -vserver vservice_name -portset portset_name -port-name lif_name
```

4. Delete the LIFs on the original nodes:

```
network interface delete -vserver vservice_name -lif lif_name
```

Unjoining the original nodes from the cluster

After the volumes have been moved to the new nodes, you unjoin the original nodes from the cluster. When you unjoin a node, the node's configuration is erased and all disks are initialized.

Steps

1. Disable high-availability configuration on the original nodes:

```
storage failover modify -node original_node_name -enabled false
```

2. Access the advanced privilege level:

```
set -privilege advanced
```

3. Identify the node that has epsilon:

```
cluster show
```

In the following example, "node0" currently holds epsilon:

```
cluster::*>
Node      Health  Eligibility  Epsilon
-----
node0     true    true         true
node1     true    true         false
node2     true    true         false
node3     true    true         false
```

4. If one of the original nodes holds epsilon, move epsilon to a different node:

- a. Remove epsilon from the original node:

```
cluster modify -node original_node_name -epsilon false
```

- b. Assign epsilon to a different node:

```
cluster modify -node new_node_name -epsilon true
```

5. From a node that will remain in the cluster, unjoin each original node from the cluster (advanced privilege level):

```
cluster unjoin -node original_node_name
```

The system displays a message similar to the following:

```
Warning: This command will unjoin node node_name from the cluster. You
must unjoin the failover partner as well. After the node is
successfully unjoined, erase its configuration and initialize all
disks by using the "Clean configuration and initialize all disks (4)"
option from the boot menu.
Do you want to continue? {y|n}: y
```

6. Enter **y** to continue.

The unjoined node is automatically rebooted and stops at the boot menu.

7. From the unjoined node's boot menu, select option **(4) Clean configuration and initialize all disks** to erase the node's configuration and initialize all disks.

The system displays a message similar to the following:

```
zero disks, reset config and install a new file system?:  
This will erase all the data on the disks, are you sure?:
```

8. Enter **y** at both prompts.
9. If the cluster has only two nodes remaining, configure high availability for the two-node cluster:

```
cluster ha modify -configured true
```

Completing the upgrade

To complete the procedure of upgrading by moving volumes, you need to configure the Service Processor (SP), install new licenses, and set up AutoSupport. You might also need to set up Storage or Volume Encryption and configure the FC or NCA ports.

Steps

1. Configure the SP on the new nodes as needed by using the `system service-processor network modify` command.
2. Install new licenses on the new nodes as needed by using the `system license add` command.
[NetApp Knowledgebase Answer 1002749: Data ONTAP 8.2 and 8.3 Licensing Overview and References](#)
3. Set up AutoSupport on the new nodes by using the `system node autosupport modify` command.
4. From each new node, send a post-upgrade AutoSupport message to technical support:

```
system node autosupport invoke -node node_name -type all -message "node_name  
successfully upgraded from platform_old to platform_new"
```
5. Restore Storage or Volume Encryption functionality by using the appropriate procedure in the *NetApp Encryption Power Guide*.

[ONTAP 9 NetApp Encryption Power Guide](#)

Use one of the following procedures, depending on whether you are using onboard or external key management:

- "Restoring onboard key management encryption keys"
 - "Restoring external key management encryption keys"
6. If the new nodes have FC ports (onboard or on FC adapters), onboard CNA ports, or a CNA card, configure the FC or CNA ports, enter the following command from the storage system prompt:

```
system node hardware unified-connect modify -node node-name -adapter adapter-name -mode  
{fc|cna} -type {target|initiator}
```

[ONTAP 9 SAN Administration Guide](#)

[Find the SAN Administration Guide for your version of Data ONTAP 8](#)

You can modify the CNA configuration only when the CNA adapters are offline.

7. Set up a switchless cluster on the new nodes if necessary.

[Migrating to a two-node switched cluster with Cisco cluster switches](#)

[Migrating to a two-node switched cluster with NetApp CN1610 cluster switches](#)

8. As needed, decommission the original systems through the NetApp Support Site to inform NetApp that the systems are no longer in operation and can be removed from support databases:
 - a. Log in to the [NetApp Support](#) site.
 - b. Click the link **My Installed Systems**.
 - c. On the **Installed Systems** page, enter the serial number of the old system in the form and then click **Go!**
 - d. On the **Decommission Form** page, fill out the form and click **Submit**.

Copyright, trademark, and machine translation

Copyright

Copyright © 2021 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

Machine translation

See important information about localized content at [netapp.com](https://www.netapp.com).

<https://www.netapp.com/company/legal/machine-translation/>