

Controller Hardware Upgrade Express Guide

For AFF and FAS Systems

March 2019 | 215-11429_AO_ur006
doccomments@netapp.com

Contents

Deciding whether to use this guide	4
Considerations for upgrading controller hardware	5
Methods for upgrading controller hardware	7
Workflow for upgrading by moving storage	8
Workflow for upgrading by moving volumes	10
Upgrading controller hardware	12
Upgrading controller hardware by moving storage	12
Preparing for the upgrade when moving storage	13
Installing the new nodes	16
Setting up the new nodes	16
Optional: Moving internal storage or converting the system to a drive shelf	18
Reassigning disks to the new nodes	20
Restoring the root volume configuration	21
Mapping network ports	22
Completing the upgrade	23
Upgrading controller hardware by moving volumes	25
Preparing for the upgrade when moving volumes	25
Installing the new nodes and joining them to the cluster	26
Creating an aggregate and moving volumes to the new nodes	26
Moving non-SAN data LIFs and cluster management LIFs to the new nodes	28
Deleting SAN LIFs from the original nodes	29
Unjoining the original nodes from the cluster	30
Completing the upgrade	31
Copyright information	33
Trademark information	34
How to send comments about documentation and receive update notifications	35
Index	36

Deciding whether to use this guide

This guide describes how to upgrade controller hardware of an All Flash FAS (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:

- You are not upgrading a MetroCluster configuration.
If you are upgrading a MetroCluster configuration, do not use this guide. You must use a nondisruptive upgrade procedure.
- Your original and new nodes are compatible and supported.
- You are upgrading a pair of nodes running ONTAP 8.3 or later to a new pair of nodes running the same release.
- You are reusing the IP addresses, netmasks, 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 replacing an individual component, see the field-replaceable unit (FRU) flyer for that component.

This guide does not describe how to use aggregate relocation to upgrade controller hardware.

Related concepts

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

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

Related information

[Using aggregate relocation to upgrade controller hardware on a pair of nodes running ONTAP 9.x](#)
[NetApp Hardware Universe](#)

Considerations for upgrading controller hardware

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

- Both the original nodes and the new nodes must be running the same and supported software version before the upgrade.
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 from the cluster. You must be sure that the number of controllers in the cluster does not exceed the supported maximum cluster size during the procedure.
NetApp Hardware Universe
- When combining different platform models of controllers in a cluster, you must be aware of the 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 procedure for upgrading by moving storage is disruptive; the procedure for upgrading by moving volumes is nondisruptive.
- If you have a FAS2220 or FAS2520 system, you can upgrade by moving volumes.
- If you have a FAS22xx or FAS25xx 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.
- If you have a FAS2240 or FAS25xx system with internal storage, you can convert the system to a drive shelf and attach it to a new node in the same cluster.
Converting a FAS2240 or FAS25xx system to a drive shelf is an optional task in the workflow for upgrading by moving storage.
- If you are upgrading to an 80xx or FAS9000 system, the 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.
NetApp Hardware Universe
- If you have an 80xx or FAS9000 controller pair and one or both controllers are AFF models, both controllers must have the same All Flash Optimized personality setting.
NetApp KB Article 1015157: How to configure a FAS8000 series controller to support the All-Flash Optimized personality
- 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, you might need to perform additional steps.

[NetApp KB Article 3015148: What are the specific steps involved in FlexArray for NetApp controller upgrades/replacements?](#)

- If your cluster has SAN hosts, you should familiarize yourself with potential issues about LUN serial number changes and the steps in resolving the issues.

[NetApp KB Article 1013497: 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

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

Related tasks

[Upgrading controller hardware by moving storage](#) on page 12

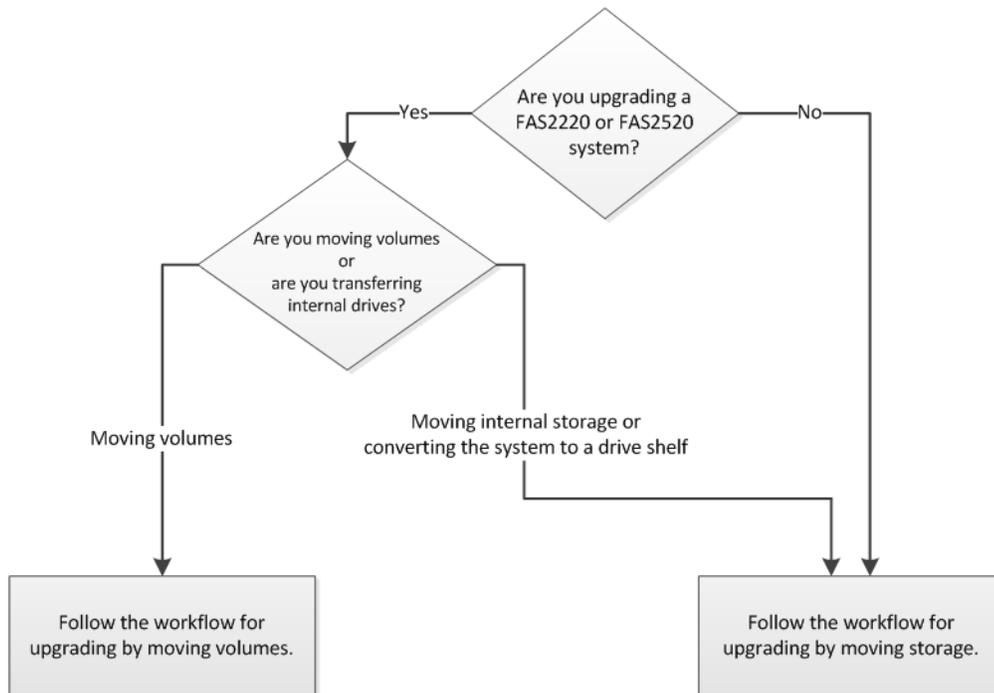
If you are upgrading by moving storage, you prepare the original nodes and set up the new nodes. Some platform models support transferring internal storage to the new nodes. You reassign disks and restore root volume configuration to the new nodes and configure network ports. Upgrading by moving storage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

[Upgrading controller hardware by moving volumes](#) on page 25

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.

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).



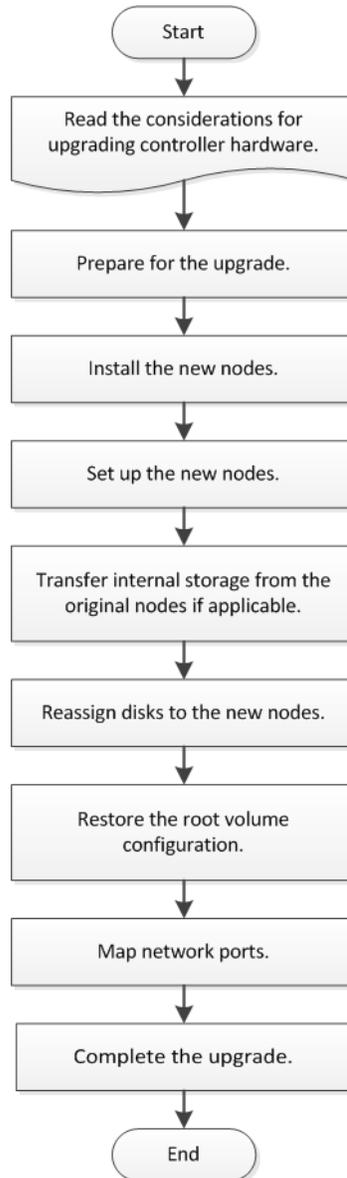
Related concepts

[Workflow for upgrading by moving storage](#) on page 8

[Workflow for upgrading by moving volumes](#) on page 10

Workflow for upgrading by moving storage

Upgrading by moving storage is a disruptive procedure. Before you start, you should read the considerations. After installing and setting up the new nodes, you can transfer storage and reassign disks to the new nodes. You then perform additional configurations and complete the upgrade.



Related concepts

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

Related tasks

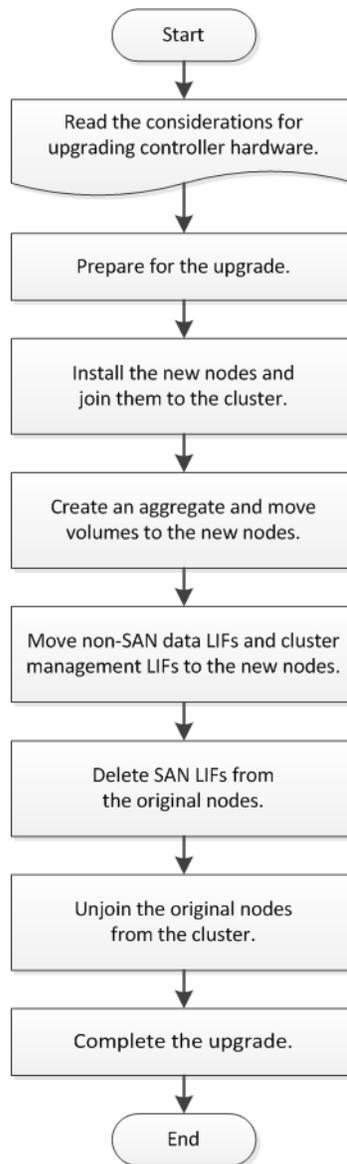
[Upgrading controller hardware by moving storage](#) on page 12

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.

Workflow for upgrading by moving volumes

Upgrading by moving volumes is a nondisruptive procedure. Before you start, you should read the considerations and prepare the original nodes. You join the new nodes to the cluster and move volumes to the new nodes. You then perform additional configurations, unjoin original nodes, reassign disks, and complete the upgrade.



Related concepts

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

Related tasks

[Upgrading controller hardware by moving volumes](#) on page 25

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.

Upgrading controller hardware

You upgrade the controller hardware either by moving the storage (a disruptive procedure) or by moving the volumes (a nondisruptive procedure).

Choices

- [Upgrading controller hardware by moving storage](#) on page 12
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 25
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.

Related concepts

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

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.

Steps

1. [Preparing for the upgrade when moving storage](#) on page 13
Before upgrading by moving storage, you must gather license information from the original nodes, plan network configuration, send an AutoSupport message about the upgrade, record the system IDs, destroy the mailboxes, power down the nodes, and remove the chassis.
2. [Installing the new nodes](#) on page 16
When you upgrade by moving storage, you begin by installing the new nodes and attaching power, console, and network connections to the new nodes.
3. [Setting up the new nodes](#) on page 16
During the process of upgrading by moving storage, you power on the new nodes, boot the software image, and configure the new nodes.
4. [Optional: Moving internal storage or converting the system to a drive shelf](#) on page 18
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.
5. [Reassigning disks to the new nodes](#) on page 20
You must reassign the disks that belonged to the original nodes to the new nodes.
6. [Restoring the root volume configuration](#) on page 21
After you install and boot the new nodes and reassign disks, you must restore configuration information from the root volume to the boot devices.
7. [Mapping network ports](#) on page 22

To enable the new node to communicate with other nodes in the cluster and with the network after the upgrade, you need to make sure that the physical ports on the new node map correctly to the physical ports on the original node.

8. [Completing the upgrade](#) on page 23

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.

Related concepts

[Workflow for upgrading by moving storage](#) on page 8

Preparing for the upgrade when moving storage

Before upgrading by moving storage, you must gather license information from the original nodes, plan network configuration, send an AutoSupport message about the upgrade, record the system IDs, destroy the mailboxes, power down the nodes, and remove the chassis.

Steps

1. Display and record license information from the original nodes by using the `system license show` command.
2. 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.
3. Record port and LIF configuration information on the original nodes:

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.
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.

To display information about...	Enter...
Broadcast domains	<code>network port broadcast-domain show</code>
IPspace information	<code>network ipspace show</code>

- 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

- As needed, adjust the configuration of the network broadcast domains on the original nodes for compatibility with that of the new nodes:

```
network port broadcast-domain modify
```

- If VLANs are configured on interface groups, remove the VLANs:

```
network port vlan delete -node node_name -port ifgrp -vlan-id VLAN_ID
```

- If any interface groups are configured on the original nodes, delete the ports that are assigned to the interface groups:

```
network port ifgrp remove-port -node node_name -ifgrp ifgrp_name -port port_name
```

- 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"
```

- Disable high availability or storage failover on each original node:

If you have a...	Enter...
Two-node cluster	<ol style="list-style-type: none"> <code>cluster ha modify -configured false</code> <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>

- Reboot the node:

```
system node reboot -node node_name
```

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

- Interrupt the reboot process by pressing Ctrl-C to display the boot menu when the system prompts you to do so.

- From the boot menu, select **(5) Maintenance mode boot** to access 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.

- 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.

Example

```
*> disk show -v
Local System ID: 118049495
DISK      OWNER          POOL      SERIAL NUMBER          HOME
-----
0a.33    node1 (118049495) Pool10    3KS6BN970000973655KL  node1 (118049495)
0a.32    node1 (118049495) Pool10    3KS6BCKD000097363ZHK  node1 (118049495)
0a.36    node1 (118049495) Pool10    3KS6BL9H000097364W74  node1 (118049495)
...
```

14. 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.

Example

```
*> 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
...
```

15. In Maintenance mode, destroy each original node's 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?
```

16. Confirm that you want to destroy the mailboxes:

```
y
```

The system displays a message 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.
*>
```

17. Exit Maintenance mode:

```
halt
```

18. Turn off the power to the original nodes, and then unplug them from the power source.

19. Label and remove all cables from the original nodes.

20. Remove the chassis containing the original nodes.

Installing the new nodes

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

Steps

1. If needed, install any adapters in the new nodes, 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 selves from the original nodes to the new nodes at this point.
3. Attach power and console connections to the new nodes, following the *Installation and Setup Instructions* for the platform.
4. Attach the network cables.
5. Transfer all remaining cables from the original node to the corresponding ports on the new node.
This includes Fibre Channel or other external disk shelf cables, and Ethernet cables.

Setting up the new nodes

During the process of upgrading by moving storage, you power on the new nodes, boot the software image, and configure the new nodes.

Steps

1. Turn on the power to the first new node, and then immediately press Ctrl-C at the console terminal to access the boot environment prompt (also called the LOADER prompt).
If you are upgrading to a system with both nodes in the same chassis, the other node also reboots. You can disregard booting the other node for now.
2. If the new node has a single-chassis configuration (with controllers in the same chassis):
 - a. Switch the console cable from the current new node to the other new node.
 - b. Turn on the power to the second new node, and then interrupt the boot process by pressing Ctrl-C at the console terminal to access the boot environment prompt.
The power should already be on if both controllers are in the same chassis.
Leave the second new node at the boot environment prompt; you return to this procedure and repeat these steps after the first new node is installed.
 - c. Switch the console cable back to the first new node.
3. At the boot environment prompt, enter the following command:

```
set-defaults
```

4. At the boot environment prompt, configure the netboot connection for a management LIF:

If IP addressing is...	Then...
DHCP	Configure the automatic connection: <code>ifconfig e0M -auto</code>

If IP addressing is...	Then...
Static	Configure the manual connection: <pre>ifconfig e0M -addr=<i>ip_addr</i> -mask=<i>netmask</i> -gw=<i>gateway</i> -dns=<i>dns_addr</i> -domain=<i>dns_domain</i></pre> <p>Note: The <code>-dns=</code> and <code>-domain=</code> arguments are optional when accessing by IP address and not by name.</p>

5. At the boot environment prompt, perform netboot on the new node:

```
netboot http://path_to_the_web-accessible_directory/netboot/kernel
```

The *path_to_the_web-accessible_directory* is the location of the downloaded `netboot.tgz` file.
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://path_to_the_web-accessible_directory/image.tgz`.

If the system prompts you for the user name, you can 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. From Maintenance mode, verify that the controller and chassis settings have the correct HA state:

```
ha-config show
```
14. If the system asks you whether to destroy the local mailboxes, enter **y** to confirm.
15. Exit Maintenance mode:

```
halt
```

The system stops at the boot environment 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 18
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 19
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.

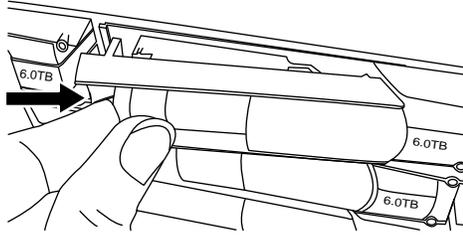
Before you begin

- 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 following illustration shows a drive with the release button located on the left of the carrier face:

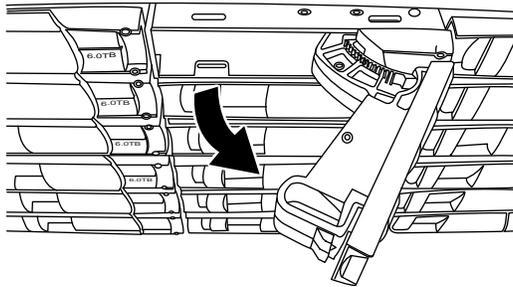


The cam handle on the carrier partially springs open, and the carrier releases from the midplane.

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.

The following illustration shows a carrier with the cam handle in its fully open position:



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 through Step 5 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.

Reassigning disks to the new nodes

You must reassign the disks that belonged to the original nodes to the new nodes.

About this task

You perform the steps in this section on both new nodes, completing each step on one node and then the other node before going on to the next step.

Steps

1. From the new node's boot environment prompt (LOADER prompt), switch to Maintenance mode:
`boot_ontap maint`
2. If the new node has a root aggregate:
 - a. Check the root aggregate of the node:
`aggr status`
 - b. Take the root aggregate of the new node offline:
`aggr offline aggr0`
 - c. Destroy the root aggregate:
`aggr destroy aggr0`
3. Display the new node's system ID:
`disk show -v`

Example

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

You should record the new node's system ID for later reference.

4. Reassign the node's spare disks, disks belonging to the root aggregate, and any SFO aggregates:
`disk reassign -s original_sysid -d new_sysid -p partner_sysID`

The *original_sysid* value is what you recorded in Step 13 on page 14, in the Preparing for an upgrade when moving storage topic.

You specify `-p partner_sysID` when shared disks are present.

Note: it is the original `partner_sysID` value that is used here, not the new value for `partner_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. Make sure that the original system's root aggregate is set to the **root** option 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

After you install and boot the new nodes and reassign disks, you must restore configuration information from the root volume to the boot devices.

About this task

You need to perform the steps in this section on both new nodes, 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 boot environment prompt (LOADER prompt):
boot_ontap
2. Press Ctrl-C to display the boot menu when the system prompts you to do so.
3. From the boot menu, select **(6) Update flash from backup config**.

The system displays the following message:

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

4. Enter **y** to continue.

The update flash process runs for several minutes, and then the system reboots. The startup process then asks you to confirm the system ID mismatch.

5. Confirm the mismatch by entering **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.

Mapping network ports

To enable the new node to communicate with other nodes in the cluster and with the network after the upgrade, you need to make sure that the physical ports on the new node map correctly to the physical ports on the original node.

Steps

1. Locate the original nodes' port and LIF configuration information that you recorded previously.

Preparing for the upgrade when moving storage on page 13

2. Record information for each new node's ports, broadcast domains, and IPspaces.

NetApp Hardware Universe

3. Make the following changes:

- a. Boot the new node to the storage system prompt if you have not already done so.

- b. Add the correct ports to the **Cluster** broadcast domain:

```
network port modify -ipspace Cluster -mtu 9000
```

Example

This example adds **Cluster** port **e1b** on “node-new”:

```
network port modify -node node-new -port e1b -ipspace Cluster -mtu
9000
```

- c. Migrate the LIFs to the new ports, once for each LIF:

```
network interface migrate -vserver vserver_name -lif lif_name -source-
node node-new -destination-node node-new -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 node-new:port
```

- f. Display the **health** state of the new node:

```
cluster show -node node-new -fields health
```

4. Delete any ports of the original node that no longer exist on the new node (advanced privilege level):

```
network port delete -node node_name -port port_name
```

5. 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 -lif lif_name -home-port port_name
```

6. Adjust the intercluster broadcast domains and migrate the intercluster LIFs, if necessary, using the same commands in Step 5.

7. Adjust any other broadcast domains and migrate the data LIFs, if necessary, using the same commands in Step 5.

8. Adjust all the LIF failover groups:

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

Example

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 “node3”:

```
network interface modify -vserver node3 -lif data1 -failover-policy broadcast-domain-wide -failover-group fg1
```

9. Display the new node's network port attributes:

```
network port show -node node-new
```

Completing the upgrade

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.

Steps

1. From the storage system prompt, display information about logical interfaces by using the `network interface show` command.
2. Delete any unused ports from the new nodes by using the `network port delete` command (advanced privilege level).

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 by using the `lun portset show` command.
 - b. Remove any unused LIFs from the port set by using the `lun portset remove` command.
4. Remove each unused LIF from the new nodes by using the `network interface delete` command.
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 by using the `system service-processor network modify` command.
7. Install new licenses on the new nodes as needed by using the `system license add` command.
[NetApp KB Article 3013749: Data ONTAP 8.2 and 8.3 Licensing Overview and References](#)
8. Set up AutoSupport on the new nodes by using the `system node autosupport modify` command.
9. 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"
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, 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.
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**.

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.

Steps

1. [Preparing for the upgrade when moving volumes](#) on page 25
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 26
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 26
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 28
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 29
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 30
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 31
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.

Related concepts

[Workflow for upgrading by moving volumes](#) on page 10

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

1. Install the new nodes and join them to the cluster:

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

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.

Example

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      voll        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
```

Example

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 vserver_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	<pre>network interface migrate -vserver vserver_name -lif lif_name -source-node source_node_name -destination-node dest_node_name -destination-port dest_port_name</pre>
All the non-SAN data LIFs and cluster-management LIFs	<pre>network interface migrate-all -node node_name</pre>

Example

The following command migrates a LIF named `datalif1` on the SVM `vs0` to the port `e0d` on `node0b`:

```
cluster::> network interface migrate -vserver vs0 -lif datalif1 -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 vserver_name -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
```

Example

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   Conn  Local      Remote      TCP Recv
Name      Name      ID    Address    Address     Size
-----
vs1       data      9     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
```

Example

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,           igroup2
        ps2       fcp     LIF3,           -
        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 vserver_name -portset portset_name -port-
name lif_name
```

4. Delete the LIFs on the original nodes:

```
network interface delete -vserver vserver_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 KB Article 3013749: 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 information

Copyright © 2018 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

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.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[*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.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

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

Index

A

- aggregates
 - creating and moving volumes to the new nodes when upgrading controller hardware [26](#)
- AutoSupport messages
 - preparing for the upgrade when moving storage [13](#)

C

- cluster management LIFs
 - moving non-SAN data LIFs and to the new nodes [28](#)
- clusters
 - considerations for upgrading controller hardware [5](#)
 - introduction to upgrading controller hardware by moving storage [12](#)
 - introduction to upgrading controller hardware by moving volumes [25](#)
 - preparing for the upgrade when moving storage [13](#)
 - unjoining the original nodes from [30](#)
 - upgrading by moving storage, converting original nodes to drive shelves [19](#)
 - workflow for upgrading by moving storage [8](#)
 - workflow for upgrading by moving volumes [10](#)
- comments
 - how to send feedback about documentation [35](#)
- considerations
 - for upgrading controller hardware [5](#)
- controller hardware
 - installing the new nodes when upgrading by moving storage [16](#)
 - introduction to upgrading by moving storage [12](#)
 - introduction to upgrading by moving volumes [25](#)
 - methods for upgrading [7](#), [12](#)
 - preparing for the upgrade when moving storage [13](#)
 - reassigning disks to the new nodes [20](#)
 - setting up the new nodes when upgrading by moving storage [16](#)
 - workflow for upgrading by moving storage [8](#)
 - workflow for upgrading by moving volumes [10](#)
- controller hardware upgrades
 - creating an aggregate and moving volumes to the new nodes [26](#)
 - installing new nodes and joining them to the cluster when moving volumes [26](#)
 - preparing for, when moving volumes [25](#)
- controller upgrades
 - deciding whether to use the Controller Hardware Upgrade Express Guide [4](#)

D

- data LIFs
 - moving non-SAN and cluster management LIFs to the new nodes [28](#)
- disks
 - reassigning to the new nodes [20](#)
- disruptive upgrade methods

- for controller hardware [7](#)
- disruptive upgrades
 - methods for upgrading controller hardware [12](#)
 - workflow for upgrading by moving storage [8](#)
- documentation
 - how to receive automatic notification of changes to [35](#)
 - how to send feedback about [35](#)
- drive shelves
 - converting original nodes to [19](#)
 - introduction to moving internal storage or converting the system to [18](#)

F

- feedback
 - how to send comments about documentation [35](#)

H

- hardware upgrades
 - deciding whether to use the Controller Hardware Upgrade Express Guide [4](#)
- hardware, controller
 - methods for upgrading [7](#)
 - preparing for the upgrade when moving storage [13](#)

I

- information
 - how to send feedback about improving documentation [35](#)
- internal drives
 - moving from an original node [18](#)
- internal storage
 - introduction to moving or converting the system to a drive shelf [18](#)

L

- license information
 - preparing for the upgrade when moving storage [13](#)
- LIFs
 - deleting SAN from the original nodes [29](#)
 - moving non-SAN data and cluster management LIFs to the new nodes [28](#)

M

- mailboxes
 - preparing for the upgrade when moving storage [13](#)
- methods
 - for upgrading controller hardware [7](#)

N

- network configurations

- preparing for the upgrade when moving storage [13](#)
- network ports
 - mapping [22](#)
- new nodes
 - installing [16](#)
 - setting up the [16](#)
- node upgrades
 - deciding whether to use the Controller Hardware Upgrade Express Guide [4](#)
- nodes
 - considerations for upgrading controller hardware [5](#)
 - creating an aggregate and moving volumes to the new nodes when upgrading controller hardware [26](#)
 - deleting SAN LIFs from the original [29](#)
 - installing and joining new nodes to the cluster when upgrading controller hardware [26](#)
 - introduction to upgrading controller hardware by moving storage [12](#)
 - introduction to upgrading controller hardware by moving volumes [25](#)
 - methods for upgrading controller hardware [12](#)
 - moving internal drives from original [18](#)
 - moving non-SAN data LIFs and cluster management LIFs to the new [28](#)
 - preparing for controller hardware upgrade when moving volumes from original [25](#)
 - preparing for the upgrade when moving storage [13](#)
 - reassigning disks to the new [20](#)
 - unjoining the original from the cluster [30](#)
 - workflow for upgrading by moving storage [8](#)
 - workflow for upgrading by moving volumes [10](#)
- non-SAN data LIFs
 - and cluster management LIFs, moving to the new nodes [28](#)
- nondisruptive upgrade methods
 - for controller hardware [7](#)
- nondisruptive upgrades
 - methods for upgrading controller hardware [12](#)
 - workflow for upgrading by moving volumes [10](#)

O

- original nodes
 - converting to drive shelves [19](#)
 - deleting SAN LIFs from [29](#)
 - moving internal drives from [18](#)
 - preparing for the upgrade when moving storage [13](#)
 - unjoining from the cluster [30](#)

P

- platforms
 - considerations for upgrading controller hardware [5](#)
 - methods for upgrading controller hardware [12](#)
- ports
 - mapping network [22](#)
- preparation
 - for the upgrade when moving storage [13](#)
- prerequisites
 - introduction to upgrading controller hardware by moving storage [12](#)

- introduction to upgrading controller hardware by moving volumes [25](#)

R

- recommendations
 - considerations for upgrading controller hardware [5](#)
- root volumes
 - restoring the configuration [21](#)

S

- SAN LIFs
 - deleting from the original nodes [29](#)
- SATA drives
 - introduction to moving internal storage or converting the system to a drive shelf [18](#)
 - moving internal drives from an original node [18](#)
- SSDs
 - introduction to moving internal storage or converting the system to a drive shelf [18](#)
 - moving internal drives from an original node [18](#)
- storage
 - installing the new nodes when upgrading by moving [16](#)
 - introduction to upgrading controller hardware by moving [12](#)
 - methods for upgrading controller hardware [7](#)
 - moving internal drives from an original node [18](#)
 - preparing for the upgrade when moving [13](#)
 - upgrade by moving, completing the upgrade [23](#)
 - upgrading by moving, mapping network ports [22](#)
 - upgrading by moving, restoring the root volume configuration [21](#)
 - upgrading by moving, setting up the new nodes [16](#)
 - workflow for upgrading by moving [8](#)
- storage controller hardware
 - considerations for upgrading [5](#)
- storage controllers
 - introduction to upgrading controller hardware by moving storage [12](#)
 - introduction to upgrading controller hardware by moving volumes [25](#)
 - methods for upgrading controller hardware [12](#)
 - preparing for the upgrade when moving storage [13](#)
 - upgrading by moving storage, converting original nodes to drive shelves [19](#)
 - workflow for upgrading by moving storage [8](#)
 - workflow for upgrading by moving volumes [10](#)
- suggestions
 - how to send feedback about documentation [35](#)
- system IDs
 - preparing for the upgrade when moving storage [13](#)

T

- Twitter
 - how to receive automatic notification of documentation changes [35](#)

U

- upgrade
 - completing the [23, 31](#)
- upgrade methods
 - for controller hardware [7](#)
- upgrades
 - methods for upgrading controller hardware [12](#)
 - preparing for, when moving storage [13](#)
 - workflow for upgrading by moving storage [8](#)
 - workflow for upgrading by moving volumes [10](#)

V

- volumes
 - creating an aggregate and moving to the new nodes when upgrading controller hardware [26](#)
 - installing new nodes and joining them to the cluster when upgrading controller hardware by moving [26](#)
 - introduction to upgrading controller hardware by moving [25](#)

- preparing for the controller hardware upgrade when moving [25](#)
- restoring the root configuration [21](#)
- upgrade by moving, completing the upgrade [31](#)
- upgrading by moving, deleting SAN LIFs from the original nodes [29](#)
- upgrading by moving, moving non-SAN data LIFs and cluster management LIFs to the new nodes [28](#)
- upgrading by moving, unjoining the original nodes from the cluster [30](#)
- workflow for upgrading by moving [10](#)

W

- workflow
 - for upgrading by moving volumes [10](#)
- workflows
 - for upgrading by moving storage [8](#)