



# **Upgrade by moving volumes or storage**

## **Upgrade controllers**

NetApp  
February 19, 2026

# Table of Contents

- Upgrade by moving volumes or storage . . . . . 1
  - Decide whether to upgrade by moving volumes or storage . . . . . 1
  - Considerations for upgrading controller hardware by moving volumes or storage . . . . . 2
    - Requirements and limitations . . . . . 2
    - Systems with internal storage . . . . . 3
    - Situations where you might need additional steps . . . . . 4
- Upgrade by moving storage . . . . . 4
  - Upgrade by moving storage workflow . . . . . 4
  - Prepare for the upgrade when moving storage . . . . . 6
  - Shut down the original nodes . . . . . 9
  - Remove ownership of disks attached to the new nodes . . . . . 11
  - Reset the default configuration on the new nodes . . . . . 13
  - Install the new nodes . . . . . 14
  - Set up the new nodes . . . . . 15
  - Optional - Move internal storage or convert to drive shelf . . . . . 17
  - Attach storage shelves and reassign disk ownership . . . . . 18
  - Restore the root volume configuration . . . . . 20
  - Complete upgrade . . . . . 22
- Upgrade by moving volumes . . . . . 33
  - Upgrade by moving volumes workflow . . . . . 33
  - Prepare for the upgrade when moving volumes . . . . . 34
  - Install the new nodes and join them to the cluster . . . . . 35
  - Move Linux iSCSI hosts to new nodes . . . . . 35
  - Create an aggregate and move volumes to the new nodes . . . . . 47
  - Move non-SAN data LIFs and cluster-management LIFs to the new nodes . . . . . 50
  - Move, delete, or create SAN LIFs . . . . . 51
  - Complete the move volumes upgrade . . . . . 53
- Upgrade AFF A250 to AFF A400 . . . . . 54
  - Convert AFF A250 to drive shelf and upgrade to AFF A400 . . . . . 54
  - Migrate LIFs and data aggregates on node2 to node1 . . . . . 55
  - Convert node2 to a drive shelf and connect to node4 . . . . . 57
  - Reassign drives from node2 to node4 . . . . . 58
  - Migrate data aggregates, epsilon, and LIFs on node1 to node4 . . . . . 60
  - Convert node1 to a drive shelf and connect to node3 . . . . . 62
  - Reassign drives from node1 to node3 . . . . . 63
  - Migrate LIFs and data aggregates on node4 to node3 . . . . . 65
- Upgrade from FAS2820 . . . . . 67
  - Upgrade from FAS2820 by converting to DS212C drive shelf . . . . . 67
  - Migrate LIFs and data aggregates on FAS2820 node2 to node1 . . . . . 68
  - Convert FAS2820 node2 to a drive shelf and connect to node4 . . . . . 71
  - Reassign drives from FAS2820 node2 to node4 . . . . . 71
  - Migrate data aggregates, epsilon, and LIFs on FAS2820 node1 to node4 . . . . . 73
  - Convert FAS2820 node1 to a drive shelf and connect to node3 . . . . . 76

Reassign drives from FAS2820 node1 to node3 .....	77
Migrate LIFs and data aggregates on FAS2820 node4 to node3 .....	78

# Upgrade by moving volumes or storage

## Decide whether to upgrade by moving volumes or storage

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

Use this procedure if you want to upgrade controller hardware in the following situations:



- ASA upgrades to an ASA r2 replacement system aren't supported. For information on migrating data from an ASA to ASA r2 system, see [Enable data access from SAN hosts to your ASA r2 storage system](#).
- When you upgrade to a system introduced in ONTAP 9.15.1 or later, ONTAP converts the storage efficiency of existing volumes and applies the new storage efficiency features that make use of the hardware offload functionality. This is an automatic background process, with no visible performance impact to the system.
  - For AFF A20, AFF A30, AFF A50, AFF A70, AFF A90, AFF A1K, AFF C30, AFF C60, and AFF C80, ONTAP converts the storage efficiency of all existing thin-provisioned volumes, including those not using storage efficiency.
  - For a FAS70 and FAS90 system, ONTAP only converts the storage efficiency of existing thin-provisioned volumes that had storage efficiency enabled before upgrading.

[Learn more about storage efficiency.](#)

- 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.
- The original and new nodes are running ONTAP 9.0 or later. NetApp recommends, when possible, that you have the same ONTAP version running on the original and new nodes.

If your controller upgrade includes mixed ONTAP versions, see [Mixed version ONTAP clusters](#) for more information.

- 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 MetroCluster configuration, see [Upgrade, refresh, or expand the MetroCluster configuration](#).

### Related information

- [Convert AFF A250 to drive shelf and upgrade to AFF A400](#) (nondisruptive procedure)
- [Upgrade from FAS2820 by converting to a drive shelf](#) (nondisruptive procedure)

- [Considerations for upgrading controller hardware](#)
- [Choose methods for upgrading controller hardware](#)
- [Where to find procedures for MetroCluster maintenance tasks](#)
- [NetApp Hardware Universe](#)

## Considerations for upgrading controller hardware by moving volumes or storage

To plan for the upgrade, you must 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

Depending on your environment, you need to consider certain factors before you start your upgrade. Get started by reviewing the table below to see the requirements and limitations you need to consider.



Before starting your controller upgrade procedure, you must review all the questions listed in the table below.

Ask yourself...	If your answer is yes, then do this..
Am I combining different controller platform models in a cluster?	<p><a href="#">Verify that you are following storage platform mixing rules.</a></p> <p>The controllers in an HA pair must be two AFF, FAS, or ASA models.</p>
Do I have different ONTAP versions running on the original and new nodes?	<ol style="list-style-type: none"> <li>1. <a href="#">Check the ONTAP versions and patch levels supported by the original and the new nodes.</a></li> <li>2. NetApp recommends, when possible, that you have the same ONTAP version running on the original and new nodes. If this is not possible, <a href="#">upgrade the ONTAP version</a> on the original nodes to the maximum supported version so that the version difference between original and new nodes is NOT greater than four. For example, ONTAP 9.8 and 9.12.1 is supported; however, ONTAP 9.8 and 9.13.1 is not supported.</li> </ol> <p><a href="#">Learn more about mixed version ONTAP clusters.</a></p>

Ask yourself...	If your answer is yes, then do this..
Do my systems contain internal drives and am I moving volumes?	<ol style="list-style-type: none"> <li>1. <a href="#">Verify that the new nodes have enough storage to accommodate storage associated with the original nodes.</a></li> <li>2. When you upgrade by moving volumes, new nodes are joined to the cluster before the original nodes are removed. You must observe the maximum cluster size.  <a href="#">Verify that the total number of controllers in the cluster during the procedure does not exceed the supported maximum cluster size.</a></li> </ol> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;">  <p>If you are upgrading an eight-node cluster serving block protocols such as FCP, iSCSI, or FCoE, verify that the new nodes are advertising the LUNs correctly. For more information, see <a href="#">SAN storage management</a>.</p> </div> <ol style="list-style-type: none"> <li>3. <a href="#">Upgrade by moving volumes</a> (a nondisruptive procedure).</li> </ol>
Am I moving internal storage or converting the system to a drive shelf?	<ol style="list-style-type: none"> <li>1. <a href="#">Verify that the current root aggregate size and the number of disks in the root aggregate either meet or exceed the specifications for the new system.</a></li> <li>2. <a href="#">Verify that the number of storage disks supported by the new system are equal to or greater than those supported by the original system.</a></li> <li>3. <a href="#">Upgrade by moving storage</a> (a disruptive procedure).</li> </ol>
Am I upgrading an HA pair in a cluster with multiple HA pairs?	Move epsilon to the node of an HA pair not undergoing a controller upgrade. For example, if you are upgrading nodeA/nodeB in a cluster with the HA pair configuration nodeA/nodeB and nodeC/nodeD, you must move epsilon to nodeC or nodeD.
Am I running ONTAP 9.6P11, 9.7P8, or later releases?	<p>NetApp recommends that you enable Connectivity, Liveliness, and Availability Monitor (CLAM) takeover to return the cluster into quorum when certain node failures occur. The <code>kernel-service</code> command requires advanced privilege level access. For more information, see the NetApp Knowledge Base article <a href="#">CLAM takeover default configuration changed</a>.</p> <p>Beginning with ONTAP 9.8, the <code>kcs-enable-takeover</code> parameter is enabled by default.</p>



You can upgrade an integrated system by moving data to new storage (moving volumes) or converting the existing integrated system into a shelf and then migrating it to a new system (moving storage). For example, you can upgrade a FAS2650 to a FAS8300 by converting the FAS2650 controller chassis into a DS224C SAS shelf and attaching it to the FAS8300. In either case, the data migration or converted shelf remains in the same switched cluster.

## Systems with internal storage

The following systems have internal storage:

Systems with internal drives			
FAS2620, FAS2650, FAS2720, and FAS2750	AFF A150, AFF A200, AFF A220, AFF A250, AFF A700s, and AFF A800	AFF C190, AFF C250, and AFF C800	ASA A150, ASA A250, ASA A800, and ASA AFF A220

- 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, you can convert the system to a drive shelf and attach it to a new node in the same cluster.



You cannot convert AFF A700s, AFF A800, AFF C800, or ASA A800 systems to a drive shelf.

- 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 you might need additional steps

- 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. See the [NetApp Hardware Universe](#).
- If your cluster has SAN hosts you might need to take steps to resolve issues with LUN serial number changes. See the Knowledge Base article [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. See the Knowledge Base article [In-Band ACP Setup and Support](#)

### Related information

- [Convert AFF A250 to drive shelf and upgrade to AFF A400](#) (nondisruptive procedure)
- [Upgrade from FAS2820 by converting to a drive shelf](#) (nondisruptive procedure)
- [Choose methods for upgrading controller hardware](#)
- [Upgrade controller hardware by moving storage](#)
- [Upgrade controller hardware by moving volumes](#)

## Upgrade by moving storage

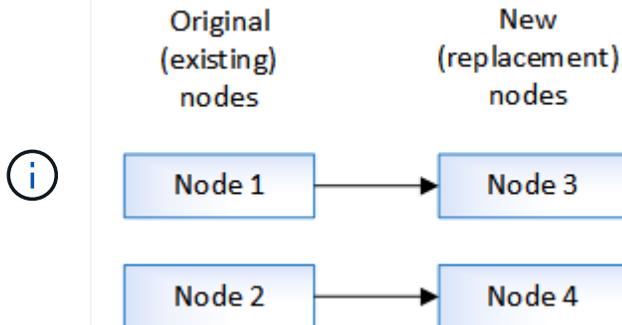
### Upgrade by moving storage workflow

Upgrading controller hardware by moving storage is a disruptive procedure. Before starting your upgrade, review the general upgrade scenarios and upgrade considerations:

- [Decide whether to upgrade by moving volumes or storage](#)
- [Considerations for upgrading controller hardware](#)

To upgrade 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.

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 these with 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.

**1**

### Prepare for the upgrade when moving storage

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

**2**

### Shut down the original nodes

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

**3**

### Remove 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. You should remove ownership of any new disks that came with node3/node4.

**4**

### Reset the default configuration on the new nodes

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

**5**

### Install 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.

**6**

### Set 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.

**7**

### Optional: Move internal storage or convert the system to a drive shelf

Optionally, if your original node is one of the supported models, you have the option to move its internal SATA drives, SSDs, or SAS drives 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**

### Attach storage shelves and reassign disk ownership

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

**9**

### Restore the root volume configuration

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

**10**

### Complete the upgrade

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

#### Related information

- [Upgrade AFF A250 to AFF A400 by converting to a drive shelf](#) (nondisruptive procedure)
- [Upgrade from FAS2820 by converting to a drive shelf](#) (nondisruptive procedure)

## Prepare 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:

```
system license show
```

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)

```
storage encryption disk show
```

- b. If any disks are associated with a non-manufacture secure ID (non-MSID) key, rekey them to an MSID key

storage encryption disk modify

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	<pre>system node run -node <i>node_name</i> sysconfig</pre>
Cluster network and node management LIFs	<pre>network interface show -role cluster,node-mgmt</pre>
Physical ports	<pre>network port show -node <i>node_name</i> -type physical</pre>
Failover groups	<pre>network interface failover-groups show -vserver <i>vserver_name</i></pre> <p>Record the names and ports of failover groups that are not clusterwide.</p>
VLAN configuration	<pre>network port vlan show -node <i>node_name</i></pre> <p>Record each network port and VLAN ID pairing.</p>
Interface group configuration	<pre>network port ifgrp show -node <i>node_name</i> -instance</pre> <p>Record the names of the interface groups and the ports assigned to them.</p>
Broadcast domains	<pre>network port broadcast-domain show</pre>
IPspace information	<pre>network ipspace show</pre>

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. Certain systems, for example, FAS8300, AFF A400, or FAS8700 systems, use ports "e0a" and "e0b" as high availability (HA) interconnect ports. If you are upgrading from a system, such as a FAS8200 or an AFF A300, to a system that uses ports "e0a" and "e0b" as HA interconnect ports, you must reassign the management and intercluster LIFs that are configured on these ports on the original system to alternative ports on the replacement system.



When ports "e0a" and "e0b" are being used as HA interconnect ports on the replacement system, any management or intercluster LIFs that are configured on these ports might lead to an upgrade failure by preventing the HA configuration from using ports "e0a" and "e0b" to boot onto the replacement system.

- a. Verify if your replacement system uses ports "e0a" and "e0b" as HA ports: [NetApp Hardware Universe](#)

- b. If required, identify the management or intercluster LIFs configured on ports "e0a" and "e0b" on your original system:

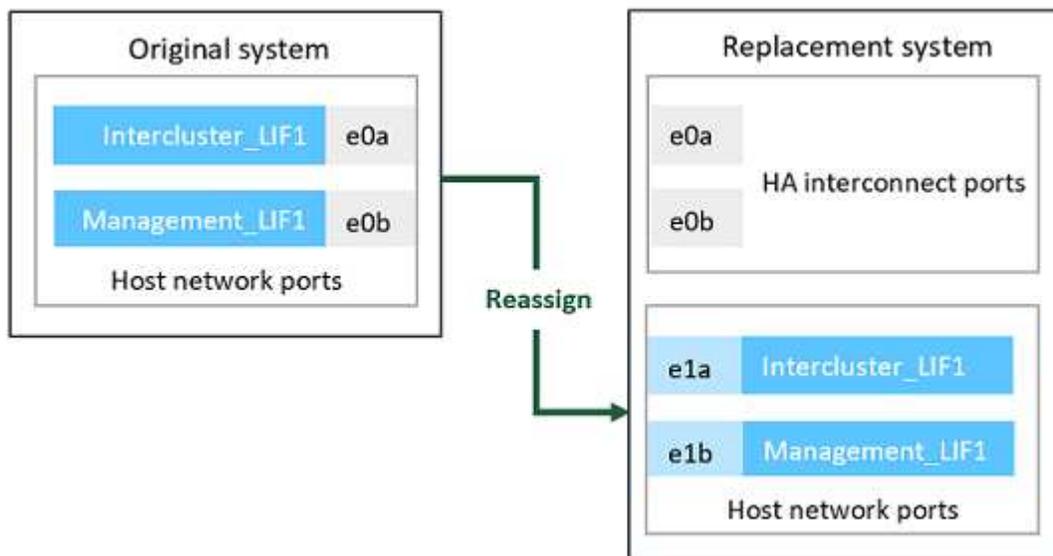
```
network interface show -home-port port_name
```

- c. If required, reassign only the affected management or intercluster LIFs to network ports that are not being used as HA ports on the replacement system:

```
network interface modify -vserver vserver_name -lif LIF_name -home-port new_port_name
```

```
network interface revert -vserver vserver_name -lif LIF_name
```

In the following example, the management and intercluster LIFs on network ports "e0a" and "e0b" are reassigned to network ports "e1a" and "e1b". Your nodes might be using different network ports because they vary by system.



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

After you install the new nodes, you might need to netboot to confirm that 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.

- Access the [NetApp Support Site](#) to download the files used for performing the netboot of the system.
- 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.
- Change to the web-accessible directory and verify that the files you need are available.

For...	Then...
<b>FAS2200, FAS2500, FAS3200, FAS6200, FAS/AFF8000 series systems</b>	<p>Extract the contents of the <code>&lt;ontap_version&gt;_image.tgz</code> file to the target directory: <code>tar -zxvf &lt;ontap_version&gt;_image.tgz</code></p> <p><b>Note:</b> 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: <code>netboot/kernel</code></p>
<b>All other systems</b>	<p>Your directory listing should contain the following file:</p> <p><code>&lt;ontap_version&gt;_image.tgz</code></p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  You do not need to extract the contents of the <code>&lt;ontap_version&gt;_image.tgz</code> file. </div>

You will use information in the directory to [set up the new nodes](#).

## Shut 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	<ol style="list-style-type: none"> <li>a. <code>cluster ha modify -configured false</code></li> <li>b. <code>storage failover modify -node <i>node_name</i> -enabled false</code></li> </ol>
Cluster with more than two nodes	<code>storage failover modify -node <i>node_name</i> -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 must be at the LOADER prompt. Use the `boot_ontap maint` command to boot into maintenance mode.

A message might appear asking you to confirm 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)  Pool10   3KS6BN970000973655KL  node1
(118049495)
0a.32    node1 (118049495)  Pool10   3KS6BCKD000097363ZHK  node1
(118049495)
0a.36    node1 (118049495)  Pool10   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.

## Remove 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.

- 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 an AFF A200 to an AFF A220 by only swapping the old AFF A200 controller module with the new AFF A220 controller module while keeping the chassis and disks of the old AFF A200 in place, you would not remove ownership of disks for the new AFF 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: 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
```

## Reset the default configuration on the new nodes

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

### About this task

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

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.

## Install 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 procedure.
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.

## Set 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 *Prepare for the upgrade when moving storage*, [Step5](#)). You need the `<ontap_version>_image.tgz` file to perform a netboot of your system.

You can also use the USB boot option to perform a netboot. See the Knowledge Base article [How to use the boot\\_recovery LOADER command for installing ONTAP for initial setup of a 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: <code>ifconfig e0M -auto</code>
Static	Configure the manual connection: <code>ifconfig e0M -addr=ip_addr -mask=netmask -gw=gateway</code>

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

For...	Then...
FAS2200, FAS2500, FAS3200, FAS6200, FAS/AFF8000 series systems	<code>netboot</code> <code>http://web_server_ip/path_to_web_accessible_directory/netboot/kernel</code>

For...	Then...
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.



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

- 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 NonDisruptive Upgrade on an HA pair". It applies to nondisruptive upgrades of software, not to upgrades of controllers.

- 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.

- 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}`
```

- 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}`
```

- 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.
- If necessary, make changes to the FC or CNA ports on the node, and then reboot the node to Maintenance mode.

### [SAN management with the CLI](#)

- You must 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.

#### 14. Exit Maintenance mode:

```
halt
```

The system stops at the LOADER prompt

## Optional - Move internal storage or convert to drive shelf

### Move 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, SSDs, or SAS drives to a drive shelf that is attached to the new node in the same cluster.



You can also choose to convert the system to a drive shelf and attach it to the new nodes.

### Before you begin

- You must have reviewed [Considerations for upgrading controller hardware](#) about moving internal drives.

Contact technical support if you need guidance specific to your configuration.

- The SATA, SSD, or SAS 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, SSD, or SAS drive carriers from the original node.

### About this task

You can move drives only within the same cluster.

### 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.



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.



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 must 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.

### Convert 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.



You can also choose to move internal drives from an original node.

#### About this task

You can move drive shelves only within the same cluster.

#### Before you begin

You must have reviewed [Considerations for upgrading controller hardware](#) about converting a node to a drive shelf. 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.

### Attach storage shelves and reassign 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 the shelves that were previously attached to node1/node2 to node3/node4.



You must not connect new shelves to node3/node4 during this upgrade procedure. You can nondisruptively connect new shelves to the system after completing the controller upgrade.

2. Verify the power supply and physical connectivity of the shelves.
3. From the node3 LOADER prompt, boot to Maintenance mode:

```
boot_ontap maint
```

4. 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.

5. 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 *Shutting down the original nodes*, [Step 5](#).
- Specify the parameter *-p partner\_sysID* only when shared disks are present.



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
```

6. 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
```

7. Enter `y` to continue.

8. Confirm 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 64-bit	root

9. Exit Maintenance mode:

```
halt
```

## Restore the root volume configuration

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



If you are performing an in-place controller upgrade to a system that uses ports "e0a" and "e0b" as high availability (HA) interconnect ports, for example, a FAS8300, an AFF A400, or a FAS8700, verify that you have [reassigned any management or intercluster LIFs](#) configured on ports "e0a" and "e0b" on the original system before starting the upgrade procedure.

### About this task

You must 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.

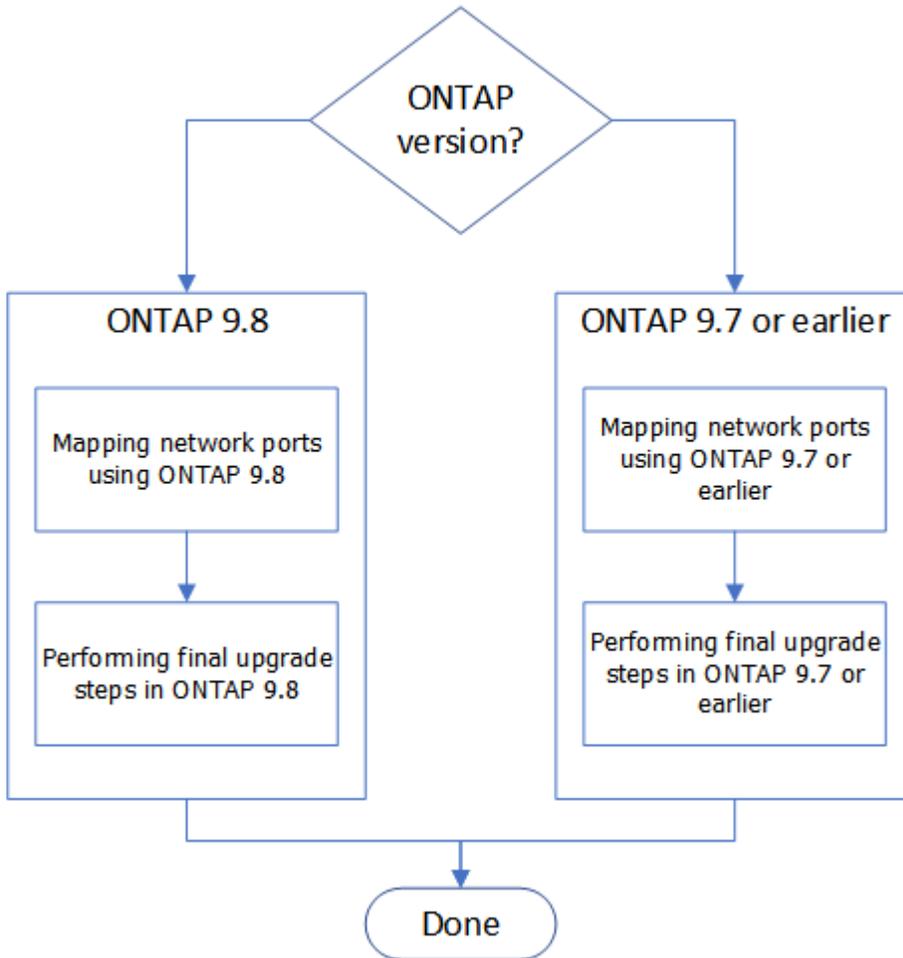
If the controller upgrade fails and reports a `rllib_port_ipspace_assign` error message, you must revert the upgrade and delete the LIFs on the network ports on the original system that are being used as HA ports on the replacement system. For more information, see [this KB article](#).

## Complete upgrade

### Complete the move storage upgrade workflow

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

You must use the procedure for your version of ONTAP.



- [Complete the upgrade in ONTAP 9.8 or later](#)
- [Complete the upgrade in ONTAP 9.7 or earlier](#)

### Complete in ONTAP 9.8 or later

#### Map network ports using ONTAP 9.8 or later

To enable node3 and node4 to communicate with each other in the cluster and with the network after the upgrade, you must confirm that 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 or later. If you are running ONTAP 9.7 or earlier, you must use the procedure in [Map network ports using ONTAP 9.7 or earlier](#).

#### About this task

You must perform these steps on node3 and node4.



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 [Map network ports using ONTAP 9.7 or earlier](#).
2. Locate the port and LIF configuration information for node1 and node2 that you recorded in *Prepare for upgrade when moving storage*, [Step 3](#).
3. Locate the information for ports, broadcast domains, and IPspaces that you recorded in *Prepare for upgrade when moving storage*, [Step 3](#).

### 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
```

- g. Depending on the ONTAP version running on the HA pair being upgraded, take one of the following actions:

If your ONTAP version is...	Then...
9.8 to 9.11.1	Verify that the cluster LIFs are listening on port 7700:  ::> network connections listening show -vserver Cluster
9.12.1 or later	Skip this step and go to <a href="#">Step 5</a> .

Port 7700 listening on cluster ports is the expected outcome as shown in the following example for a two-node cluster:

```
Cluster::> network connections listening show -vserver Cluster
Vserver Name      Interface Name:Local Port      Protocol/Service
-----
Node: NodeA
Cluster           NodeA_clus1:7700              TCP/ctlopcp
Cluster           NodeA_clus2:7700              TCP/ctlopcp
Node: NodeB
Cluster           NodeB_clus1:7700              TCP/ctlopcp
Cluster           NodeB_clus2:7700              TCP/ctlopcp
4 entries were displayed.
```

- h. For each cluster LIF that is not listening on port 7700, set the administrative status of the LIF to down and then up:

```
::> net int modify -vserver Cluster -lif cluster-lif -status-admin down; net
int modify -vserver Cluster -lif cluster-lif -status-admin up
```

Repeat substep (g) to verify that the cluster LIF is now listening on port 7700.

5. Modify the broadcast domain memberships of physical ports hosting data LIFs. You can do this manually, as shown in [Map network ports using ONTAP 9.7 or earlier, Step 7](#). NetApp recommends using the enhanced network reachability scan and repair procedure introduced in ONTAP 9.8, as shown in the following Step 5, substeps (a) to (g).

- 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 enable 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:

```
network port 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
```



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

6. If there are no VLANs configured on the system, go to [Step 7](#). 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:

```
cluster controller-replacement network 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:

```
cluster controller-replacement network 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. Beginning with 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 vserver_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](#).

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 vserver_name -status-admin down
```

- b. Check for any LIFs that are operationally down: `network interface show -vserver vserver_name -status-oper down`

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

```
network interface modify -vserver vserver_name -lif lif -home-port home_port
```



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 [Perform final upgrade steps](#)

in [ONTAP 9.8 or later](#).

### Perform final upgrade steps in ONTAP 9.8 or later

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 or later. 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](#).

### 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](#).
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: <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>

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
```

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 one of the following procedures, depending on whether you are using onboard or External Key Management:

- [Restore onboard key management encryption keys](#)
- [Restore 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}
```

### [SAN management with the CLI](#)

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. If required, re-create the non-default user accounts that you had for the baseboard management controller (BMC) on the old system:

- a. Change or reset the BMC admin user account password.

The BMC admin user account password is either blank (no password) or the same as the system admin user account password.

- b. Re-create the non-default BMC user accounts by using the `security login create` command with `application` set to "service-processor", as shown in the following example:

```
security login create -user-or-group-name bmcuser -application service-  
processor -authentication-method password -role admin
```



Admin privileges are required to create a user account in the BMC.

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.

### Complete in ONTAP 9.7 or earlier

#### Map 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 must confirm that 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 or later, you must use the procedure in [Map network ports using ONTAP 9.8 or later](#).

### About this task

You must perform these steps on node3 and node4.



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 or later, **STOP**. You must use the procedure in [Map network ports using ONTAP 9.8 or later](#).
2. Locate the port and LIF configuration information for node1 and node2 that you recorded in *Prepare for upgrade when moving storage*, [Step 3](#).
3. Locate the information for ports, broadcast domains, and IPspaces that you recorded in *Prepare for upgrade when moving storage*, [Step 3](#).

#### [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 vservice_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 -ip-space Cluster -broadcast  
-domain Cluster -ports node1:port
```

f. Display the health state of node3 and node4:

```
cluster show -node node1 -fields health
```

g. Each cluster LIF must be listening on port 7700. Verify that the cluster LIFs are listening on port 7700:

```
::> network connections listening show -vserver Cluster
```

Port 7700 listening on cluster ports is the expected outcome as shown in the following example for a two-node cluster:

```
Cluster::> network connections listening show -vserver Cluster
Vserver Name      Interface Name:Local Port      Protocol/Service
-----
Node: NodeA
Cluster           NodeA_clus1:7700              TCP/ctlopcp
Cluster           NodeA_clus2:7700              TCP/ctlopcp
Node: NodeB
Cluster           NodeB_clus1:7700              TCP/ctlopcp
Cluster           NodeB_clus2:7700              TCP/ctlopcp
4 entries were displayed.
```

h. For each cluster LIF that is not listening on port 7700, set the administrative status of the LIF to down and then up:

```
::> net int modify -vserver Cluster -lif cluster-lif -status-admin down; net  
int modify -vserver Cluster -lif cluster-lif -status-admin up
```

Repeat substep (g) to verify that the cluster LIF is now listening on port 7700.

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](#).

9. Adjust any other broadcast domains and migrate the data LIFs, if necessary, using the commands in [Step 7](#).

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 [Perform final upgrade steps in ONTAP 9.7 or earlier](#).

### Perform the final upgrade steps in ONTAP 9.7 or earlier

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 9.7 or earlier. If you are running ONTAP 9.8 or later, you must use the procedure in [Perform the final upgrade steps in ONTAP 9.8 or later](#).

### Steps

1. If your system is running ONTAP 9.8 or later, **STOP**. You must use the procedure in [Perform final upgrade steps in ONTAP 9.8 or later](#).
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
```

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 one of the following procedures, depending on whether you are using onboard or External Key Management:

- [Restore onboard key management encryption keys](#)

- [Restore 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}
```

### [SAN management with the CLI](#)

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:

- Log in to the [NetApp Support](#) site.
- Click the link **My Installed Systems**.
- On the Installed Systems page, enter the serial number of the old system in the form and then click **Go!**
- On the Decommission Form page, fill out the form and click **Submit**.

### **After you finish**

You have completed the upgrade procedure.

## **Upgrade by moving volumes**

### **Upgrade by moving volumes workflow**

Upgrading controller hardware by moving volumes is a nondisruptive procedure. Before starting your upgrade, review the general upgrade scenarios and upgrade considerations:

- [Decide whether to upgrade by moving volumes or storage](#)
- [Considerations for upgrading controller hardware](#)

To upgrade 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 remove the original nodes from the cluster.

**1**

#### **[Prepare for the upgrade when moving volumes](#)**

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

**2**

#### **[Install the new nodes and join them to the cluster](#)**

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

**3**

### Move Linux iSCSI hosts to the new nodes

Before moving iSCSI SAN volumes to new nodes, you create new iSCSI connections and rescan the iSCSI paths to the new nodes.

**4**

### Create an aggregate and move volumes to the new nodes

You create at least one 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

**5**

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

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

**6**

### Move, delete, or create SAN LIFS

Depending on your cluster contents and cluster environment, you move, delete, or create SAN LIFs, or re-create deleted SAN LIFs.

**7**

### Remove the original nodes from the cluster

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

**8**

### Complete the upgrade

To complete the procedure of upgrading by moving volumes, you 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.

## Prepare for the upgrade when moving volumes

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

### Steps

1. Display the volumes on the original nodes:

```
volume show
```

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:

```
system license show
```

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):

```
storage encryption disk show
```

- b. If any disks are associated with a non-manufacture secure ID (non-MSID) key, rekey them to an MSID key:

```
storage encryption disk modify
```

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"
```

### What's next?

[Install the new nodes and join them to the cluster](#)

## Install the new nodes and join 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	<a href="#">Cluster expansion administration</a>
Releases before ONTAP 9.0	<a href="#">Find the Cluster Expansion Express Guide for your version of Data ONTAP 8</a>

### What's next?

[Move Linux iSCSI hosts to new nodes](#)

## Move Linux iSCSI hosts to new nodes

Before moving iSCSI SAN volumes to new nodes, you must create new iSCSI

connections and rescan the iSCSI paths to the new nodes.

If you do not need to move iSCSI SAN volumes when upgrading by moving volumes, you can skip this procedure and go to [Create an aggregate and move volumes to the new nodes](#).

#### About this task

- IPv4 interfaces are created when you set up the new iSCSI connections.
- The host commands and examples are specific to Linux operating systems.

#### Step 1: Set up new iSCSI connections

To move the iSCSI connections, you set up new iSCSI connections to the new nodes.

#### Steps

1. Create iSCSI interfaces on the new nodes and check ping connectivity from the iSCSI hosts to the new interfaces on the new nodes.

##### [Create network interfaces](#)

All iSCSI interfaces from the SVM should be reachable by the iSCSI host.

2. On the iSCSI host, identify the existing iSCSI connections from the host to the old node:

```
iscsiadm -m session
```

```
[root@scspr1789621001 ~]# iscsiadm -m session
tcp: [1] 10.230.68.236:3260,1156 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-flash)
tcp: [2] 10.230.68.237:3260,1158 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-flash)
```

3. On the new node, verify the connections from the new node:

```
iscsi session show -vserver <svm-name>
```

```

node_A_1-new::*> iscsi session show -vserver vsa_1
  Tpgroup Initiator Initiator
Vserver Name TSIH Name ISID Alias
-----
-----
vsa_1 iscsi_lf_n1_p1_4 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:01
scspr1789621001.gdl.englab.netapp.com
vsa_1 iscsi_lf_n2_p1_4 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:02
scspr1789621001.gdl.englab.netapp.com
2 entries were displayed.

```

4. On the new node, list the iSCSI interfaces in ONTAP for the SVM that contains the interfaces:

```
iscsi interface show -vserver <svm-name>
```

```

sti8200mcchtp001htp_siteA::*> iscsi interface show -vserver vsa_1
  Logical Status Curr Curr
Vserver Interface TPGT Admin/Oper IP Address Node Port Enabled
-----
-----
vsa_1 iscsi_lf_n1_p1_1156 up/up 10.230.68.236 sti8200mcc-htp-001 e0g
true
vsa_1 iscsi_lf_n1_p2_1157 up/up fd20:8b1e:b255:805e::78c9 sti8200mcc-
htp-001 e0h true
vsa_1 iscsi_lf_n2_p1_1158 up/up 10.230.68.237 sti8200mcc-htp-002 e0g
true
vsa_1 iscsi_lf_n2_p2_1159 up/up fd20:8b1e:b255:805e::78ca sti8200mcc-
htp-002 e0h true
vsa_1 iscsi_lf_n3_p1_1183 up/up 10.226.43.134 sti8200mccip-htp-005 e0c
true
vsa_1 iscsi_lf_n4_p1_1188 up/up 10.226.43.142 sti8200mccip-htp-006 e0c
true
6 entries were displayed.

```

5. On the iSCSI host, run discovery on any one of the iSCSI IP addresses on the SVM to discover the new targets:

```
iscsiadm -m discovery -t sendtargets -p iscsi-ip-address
```

Discovery can be run on any IP address of the SVM, including non-iSCSI interfaces.

```
[root@scspr1789621001 ~]# iscsiadm -m discovery -t sendtargets -p
10.230.68.236:3260
10.230.68.236:3260,1156 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
10.226.43.142:3260,1188 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
10.226.43.134:3260,1183 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
10.230.68.237:3260,1158 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
```

6. On the iSCSI host, login to all the discovered addresses:

```
iscsiadm -m node -L all -T node-address -p portal-address -l
```

```
[root@scspr1789621001 ~]# iscsiadm -m node -L all -T iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 -p
10.230.68.236:3260 -l
Logging in to [iface: default, target: iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6, portal:
10.226.43.142,3260] (multiple)
Logging in to [iface: default, target: iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6, portal:
10.226.43.134,3260] (multiple)
Login to [iface: default, target: iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6, portal:
10.226.43.142,3260] successful.
Login to [iface: default, target: iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6, portal:
10.226.43.134,3260] successful.
```

7. On the iSCSI host, verify the login and connections:

```
iscsiadm -m session
```

```
[root@scspr1789621001 ~]# iscsiadm -m session
tcp: [1] 10.230.68.236:3260,1156 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-flash)
tcp: [2] 10.230.68.237:3260,1158 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-flash)
tcp: [3] 10.226.43.142:3260,1188 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-flash)
```

8. On the new node, verify the login and connection with the host:

```
iscsi initiator show -vserver <svm-name>
```

```
sti8200mcchtp001htp_siteA::*> iscsi initiator show -vserver vsa_1
  Tpgroup Initiator
Vserver Name          TSIH Name          ISID
Igroup Name
-----
vsa_1 iscsi_lf__n1_p1_ 4 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:01 igroup_linux
vsa_1 iscsi_lf__n2_p1_ 4 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:02 igroup_linux
vsa_1 iscsi_lf__n3_p1_ 1 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:04 igroup_linux
vsa_1 iscsi_lf__n4_p1_ 1 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:03 igroup_linux
4 entries were displayed.
```

## Result

At the end of this task, the host can see all iSCSI interfaces (on the old and new nodes) and is logged in to all those interfaces.

LUNs and volumes are still physically hosted on the old nodes. Because LUNs are reported only on the old node interfaces, the host will show only paths over the old nodes. To see this, run the `sanlun lun show -p` and `multipath -ll -d` commands on the host and examine the command outputs.

```

[root@scspr1789621001 ~]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host vserver
path path /dev/ host vserver
state type node adapter LIF
-----
up primary sdk host3 iscsi_lf__n2_p1_
up secondary sdh host2 iscsi_lf__n1_p1_
[root@scspr1789621001 ~]# multipath -ll -d
3600a098038304646513f4f674e52774b dm-5 NETAPP ,LUN C-Mode
size=2.0G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|+- policy='service-time 0' prio=50 status=active
| `-- 3:0:0:4 sdk 8:160 active ready running
`+- policy='service-time 0' prio=10 status=enabled
`-- 2:0:0:4 sdh 8:112 active ready running

```

## Step 2: Add the new nodes as reporting nodes

After setting up the connections to the new nodes, you add the new nodes as the reporting nodes.

### Steps

1. On the new node, list reporting nodes for LUNs on the SVM:

```

lun mapping show -vserver <svm-name> -fields reporting-nodes -ostype
linux

```

The following reporting nodes are local nodes as LUNs are physically on old nodes node\_A\_1-old and node\_A\_2-old.

```

node_A_1-new::*> lun mapping show -vserver vsa_1 -fields reporting-nodes
-ostype linux
vserver path                                igroup      reporting-nodes
-----
-----
vsa_1    /vol/vsa_1_vol1/lun_linux_2  igroup_linux node_A_1-old,node_A_2-
old
.
.
.
vsa_1    /vol/vsa_1_vol9/lun_linux_19 igroup_linux node_A_1-old,node_A_2-
old
12 entries were displayed.

```

2. On the new node, add reporting nodes:

```

lun mapping add-reporting-nodes -vserver <svm-name> -path
/vol/vsa_1_vol*/lun_linux_* -nodes node1,node2 -igroup <igroup_name>

```

```

node_A_1-new::*> lun mapping add-reporting-nodes -vserver vsa_1 -path
/vol/vsa_1_vol*/lun_linux_* -nodes node_A_1-new,node_A_2-new
-igroup igroup_linux
12 entries were acted on.

```

3. On the new node, verify that the newly added nodes are present:

```

lun mapping show -vserver <svm-name> -fields reporting-nodes -ostype
linux vserver path igroup reporting-nodes

```

```

node_A_1-new::*> lun mapping show -vserver vsa_1 -fields reporting-nodes
-ostype linux vserver path igroup reporting-nodes
-----
-----
-----
vsa_1 /vol/vsa_1_voll/lun_linux_2 igroup_linux node_A_1-old,node_A_2-
old,node_A_1-new,node_A_2-new
vsa_1 /vol/vsa_1_voll/lun_linux_3 igroup_linux node_A_1-old,node_A_2-
old,node_A_1-new,node_A_2-new
.
.
.
12 entries were displayed.

```

4. The `sg3-utils` package must be installed on the Linux host. This prevents a `rescan-scsi-bus.sh` utility not found error when you rescan the Linux host for the newly mapped LUNs using the `rescan-scsi-bus` command.

On the host, verify that the `sg3-utils` package is installed:

- For a Debian based distribution:

```
dpkg -l | grep sg3-utils
```

- For a Red Hat based distribution:

```
rpm -qa | grep sg3-utils
```

If required, install the `sg3-utils` package on the Linux host:

```
sudo apt-get install sg3-utils
```

5. On the host, rescan the SCSI bus on the host and discover the newly added paths:

```
/usr/bin/rescan-scsi-bus.sh -a
```

```

[root@stemgr]# /usr/bin/rescan-scsi-bus.sh -a
Scanning SCSI subsystem for new devices
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
  Scanning for device 2 0 0 0 ...
.
.
.
OLD: Host: scsi5 Channel: 00 Id: 00 Lun: 09
  Vendor: NETAPP Model: LUN C-Mode Rev: 9800
  Type: Direct-Access ANSI SCSI revision: 05
0 new or changed device(s) found.
0 remapped or resized device(s) found.
0 device(s) removed.

```

6. On the iSCSI host, list the newly added paths:

```
sanlun lun show -p
```

Four paths are shown for each LUN.

```

[root@stemgr]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host vserver
path path /dev/ host vserver
state  type      node   adapter  LIF
-----
up     primary   sdk    host3     iscsi_lf__n2_p1_
up     secondary sdh     host2     iscsi_lf__n1_p1_
up     secondary sdag    host4     iscsi_lf__n4_p1_
up     secondary sdah    host5     iscsi_lf__n3_p1_

```

7. On the new node, move the volume/volumes containing LUNs from the old nodes to the new nodes.

```

node_A_1-new::*> vol move start -vserver vsa_1 -volume vsa_1_vol1
-destination-aggregate sti8200mccip_htp_005_aggr1
[Job 1877] Job is queued: Move "vsa_1_vol1" in Vserver "vsa_1" to
aggregate "sti8200mccip_htp_005_aggr1". Use the "volume move show
-vserver
vsa_1 -volume vsa_1_vol1" command to view the status of this operation.
node_A_1-new::*> vol move show
Vserver   Volume           State           Move           Phase           Percent-
Complete  Time-To-Complete
-----
-----
vsa_1     vsa_1_vol1       healthy         -              initializing    -
-

```

8. When the volume move to the new nodes is complete, verify that the volume is online:

```

volume show -state

```

9. The iSCSI interfaces on the new nodes where the LUN now resides are updated as primary paths. If the primary path is not updated after the volume move, run `/usr/bin/rescan-scsi-bus.sh -a` and `multipath -v3` on the host or simply wait for multipath rescanning to take place.

In the following example, the primary path is a LIF on the new node.

```

[root@stemgr]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host vserver
path path /dev/ host vserver
state type node adapter LIF
-----
up primary sdag host4 iscsi_lf__n4_p1_
up secondary sdk host3 iscsi_lf__n2_p1_
up secondary sdh host2 iscsi_lf__n1_p1_
up secondary sdah host5 iscsi_lf__n3_p1_

```

### Step 3: Remove reporting nodes and rescan paths

You must remove the reporting nodes and rescan the paths.

#### Steps

1. On the new node, remove remote reporting nodes (the new nodes) for the Linux LUNs:

```
lun mapping remove-reporting-nodes -vserver <svm-name> -path * -igroup  
<igroup_name> -remote-nodes true
```

In this case, the remote nodes are old nodes.

```
node_A_1-new::*> lun mapping remove-reporting-nodes -vserver vsa_1 -path  
* -igroup igroup_linux -remote-nodes true  
12 entries were acted on.
```

2. On the new node, check reporting nodes for the LUNs:

```
lun mapping show -vserver <svm-name> -fields reporting-nodes -ostype  
linux
```

```
node_A_1-new::*> lun mapping show -vserver vsa_1 -fields reporting-nodes  
-ostype linux  
vserver  path                                     igroup      reporting-nodes  
-----  -----  
-----  
vsa_1    /vol/vsa_1_vol1/lun_linux_2  igroup_linux node_A_1-  
new,node_A_2-new  
vsa_1    /vol/vsa_1_vol1/lun_linux_3  igroup_linux node_A_1-  
new,node_A_2-new  
vsa_1    /vol/vsa_1_vol2/lun_linux_4  group_linux  node_A_1-  
new,node_A_2-new  
.  
.  
.  
12 entries were displayed.
```

3. The `sg3-utils` package must be installed on the Linux host. This prevents a `rescan-scsi-bus.sh` utility not found error when you rescan the Linux host for the newly mapped LUNs using the `rescan-scsi-bus` command.

On the host, verify that the `sg3-utils` package is installed:

- For a Debian based distribution:

```
dpkg -l | grep sg3-utils
```

- For a Red Hat based distribution:

```
rpm -qa | grep sg3-utils
```

If required, install the `sg3-utils` package on the Linux host:

```
sudo apt-get install sg3-utils
```

4. On the iSCSI host, rescan the SCSI bus:

```
/usr/bin/rescan-scsi-bus.sh -r
```

The paths that are removed are the paths from the old nodes.

```

[root@scspr1789621001 ~]# /usr/bin/rescan-scsi-bus.sh -r
Syncing file systems
Scanning SCSI subsystem for new devices and remove devices that have
disappeared
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
sg0 changed: LU not available (PQual 1)
REM: Host: scsi2 Channel: 00 Id: 00 Lun: 00
DEL: Vendor: NETAPP Model: LUN C-Mode Rev: 9800
Type: Direct-Access ANSI SCSI revision: 05
sg2 changed: LU not available (PQual 1)
.
.
.
OLD: Host: scsi5 Channel: 00 Id: 00 Lun: 09
Vendor: NETAPP Model: LUN C-Mode Rev: 9800
Type: Direct-Access ANSI SCSI revision: 05
0 new or changed device(s) found.
0 remapped or resized device(s) found.
24 device(s) removed.
[2:0:0:0]
[2:0:0:1]
.
.
.

```

5. On the iSCSI host, verify that only paths from the new nodes are visible:

```
sanlun lun show -p
```

```
multipath -ll -d
```

### What's next?

[Create an aggregate and move volumes to the new nodes](#)

## Create an aggregate and move 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 the required data protection procedure.](#)

- If you are moving iSCSI SAN volumes, verify that you have [created new iSCSI connections](#).



For each storage virtual machine (SVM), NetApp recommends moving all the non-root volumes in the cluster before moving the root volume, and performing this procedure on one SVM at a time.

## 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 storage virtual machine (SVM) as the aggregate on the original node from which you want to move the volumes:

```
vserver add-aggregates
```

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:

```
vserver show -vserver svm_name
```

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       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
```

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.



If your storage configuration includes volumes with encryption enabled, follow the steps in [Enable encryption on an existing volume with the volume move start command](#) to move these volumes.

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
```

## What's next?

[Move non-SAN data LIFs and cluster-management LIFs to the new nodes](#)

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

After you have moved the volumes from the original nodes, you must 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. Log in with the cluster-management LIF, and list all LIFs on the original nodes (comma-separated list):

```
network interface show -curr-node <list_of_original_node_names>
```

2. 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|lifgrp>}
```

3. Take one of the following actions:

If you want to migrate...	Then enter...
A specific LIF	<pre>network interface migrate -vserver &lt;vserver_name&gt; -lif &lt;lif_name&gt; -destination -node &lt;dest_node_name&gt; -destination-port &lt;dest_port_name&gt;</pre>
All the non-SAN data LIFs and cluster-management LIFs	<pre>network interface migrate-all -node &lt;node_name&gt;</pre>

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
```

4. 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
```

5. 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>}
```

### What's next?

[Move, delete, or create SAN LIFs](#)

## Move, delete, or create SAN LIFs

Depending on your cluster contents and cluster environment, you must move, delete, or create SAN LIFs, or re-create deleted SAN LIFs.

### Considerations for moving SAN LIFs

You only need to move the SAN LIFs if you are changing the contents of your cluster, for example, by adding nodes to the cluster or deleting nodes from the cluster. When you move a LIF, you do not have to re-zone your FC fabric or create new iSCSI sessions between the attached hosts of your cluster and the new target interface.

You can move a SAN LIF by using the `network interface modify` command. To move a SAN LIF, you must take the LIF offline, move the LIF to a different home node or port, and then bring it back online in its new location. Asymmetric Logical Unit Access (ALUA) provides redundant paths and automatic path selection as part of any ONTAP SAN solution. Therefore, when the LIF is taken offline for the movement, there is no I/O interruption. The host simply retries and then moves I/O to another LIF.

During the LIF movements, you can nondisruptively perform the following tasks:

- Replace one HA pair of a cluster with an upgraded HA pair in a way that is transparent to the hosts accessing LUN data
- Upgrade a target interface card
- Shift the resources of a storage virtual machine (SVM) from one set of nodes in a cluster to another set of nodes in the same cluster
- When the host server is online, you can move a SAN LUN to a new HA pair without disrupting the host server access to the LUN data

For more information, see the [SAN LIF movement](#) procedure in the *SAN storage management* documentation.

## Delete SAN LIFs no longer required

If the cluster is in a SAN environment, you must delete any SAN LIFs that you no longer require from the original nodes before you can remove 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
          Tpgroup      Conn  Local          Remote          TCP
Recv
Vserver  Name      TSIH  ID   Address          Address          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,
          LIF2                igroup1
          ps1          iscsi    LIF3                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 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
```

### Create new SAN LIFs or re-create deleted SAN LIFs

Depending on your cluster environment requirements, you might decide to create new SAN LIFs or re-create SAN LIFs that you deleted earlier in this procedure. You can create or re-create SAN LIFs by using the [network interfaces creation](#) procedure in the *Cluster Management Using OnCommand® System Manager* documentation.

### What's next?

[Remove the original nodes from the cluster](#)

## Complete the move volumes upgrade

To complete the procedure of upgrading by moving volumes, you must 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.

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

```
system service-processor network modify
```

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

```
system license add
```

3. Set up AutoSupport on the new nodes:

```
system node autosupport modify
```

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 one of the following procedures, depending on whether you are using onboard or External Key Management:

- [Restore onboard key management encryption keys](#)
- [Restore 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 by entering 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}
```

### [SAN management with the CLI](#)

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**.

## Upgrade AFF A250 to AFF A400

### Convert AFF A250 to drive shelf and upgrade to AFF A400

Perform a nondisruptive upgrade from an AFF A250 system to an AFF A400 system by converting the AFF A250 nodes to NS224 drive shelves. Then, connect them to the AFF A400 replacement nodes. This moves the AFF A250 onboard storage to the AFF A400 replacement system.

#### Before you begin

Review the general upgrade scenarios and upgrade considerations for upgrading by moving volumes or storage:

- [Decide whether to upgrade by moving volumes or storage](#)
- [Considerations for upgrading controller hardware](#)

#### About this task

The AFF A250 high-availability (HA) pair controllers are node1 and node2 and the replacement AFF A400 HA pair controllers are node3 and node4.

1

#### Migrate LIFs and data aggregates on node2 to node1

Before converting AFF A250 node2 to a drive shelf, migrate the logical interfaces (LIFs) and data aggregates on node2 to node1.

2

#### Convert node2 to a drive shelf and connect to node4

Convert AFF A250 node2 to an NS224 drive shelf and then connect to AFF A400 node4 before reassigning drives from node2 to node4.

3

#### Reassign drives from node2 to node4

After converting AFF A250 node2 to an NS224 drive shelf and connecting to AFF A400 node4, reassign the

drives that previously belonged to node2 to node4.

4

#### Migrate data aggregates, epsilon, and LIFs on node1 to node4

Before converting AFF A250 node1 to a drive shelf, migrate the data aggregates, epsilon, and LIFs on node1 to AFF A400 node4.

5

#### Convert node1 to a drive shelf and connect to node3

Convert AFF A250 node1 to an NS224 drive shelf and then connect to AFF A400 node3 before reassigning drives from node1 to node3.

6

#### Reassign drives from node1 to node3

After converting AFF A250 node1 to an NS224 drive shelf and connecting to AFF A400 node3, reassign the drives that previously belonged to node1 to node3.

7

#### Migrate LIFs and data aggregates on node4 to node3

To complete the upgrade, connect node3 to node4 and then migrate the data LIFs and data aggregates on node4 to node3.

## Migrate LIFs and data aggregates on node2 to node1

Before converting AFF A250 node2 to a drive shelf, you migrate the logical interfaces (LIFs) and data aggregates on node2 to node1.

### Before you begin

Verify that you meet the following requirements:

- The AFF A250 and AFF A400 controllers are running the same ONTAP release and patch version.
  - You must netboot and install the ONTAP version on each AFF A400 that is identical to the version running on the AFF A250.
  - Both the primary and backup boot image of each AFF A400 must have the same ONTAP version.
  - If the AFF A400 clusters were previously configured, you must clear any residual cluster configuration by performing a `wipeconfig` from the boot menu.
- Both AFF A400 controllers are on standby at the LOADER prompt.
- You have all the appropriate cabling on-hand.



### About this task

The following steps are performed on AFF A250 node1.

### Steps

1. Access the advanced privilege level:

```
set -privilege advanced
```

2. Disable storage failover automatic giveback:

```
storage failover modify -node node1 -auto-giveback false
```

3. Disable auto-revert of the LIFs across both nodes of the HA pair:

```
network interface modify -lif * -auto-revert false
```

4. Display the status of all data network LIFs:

```
network interface show -role data
```

5. Display the status of the cluster management LIFs:

```
network interface show -role cluster_mgmt
```

6. Migrate all data LIFs from the storage virtual machines hosted on node2:

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



This command only migrates non-SAN LIFs. You cannot use it to migrate iSCSI and FCP LIFs.

7. Display the status of all data LIFs in the cluster:

```
network interface show -role data
```

8. If any LIFs are down, set the administrative status of the LIFs to up by entering the following command, once for each LIF:

```
network interface modify -vserver vserver_name -lif lif_name -status-admin up
```

9. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

10. Display failover eligibility:

```
storage failover show
```

11. Migrate the data aggregates on node2 to node1:

```
storage aggregate relocation start -aggregate aggregate_name -node node2  
-destination node1
```

12. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

13. Display the status of all data volumes in the cluster:

```
volume show
```

14. Display the ha status and ownership of epsilon:

```
cluster show
```

15. Disable cluster ha:

```
cluster ha modify -configured false
```

16. Display the ha status and ownership of epsilon:

```
cluster show
```

17. Halt node2:

```
halt -node node2 -inhibit-takeover true -ignore-quorum-warnings true
```

### What's next?

[Convert node2 to a drive shelf and connect to node4](#)

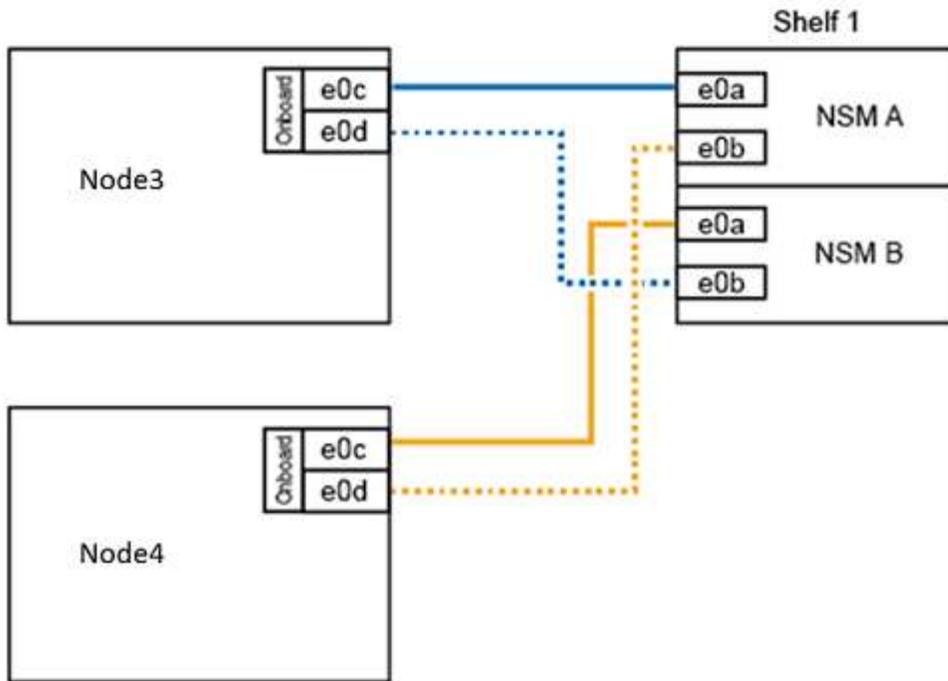
## Convert node2 to a drive shelf and connect to node4

Convert AFF A250 node2 to an NS224 drive shelf and then connect to AFF A400 node4 before reassigning drives from node2 to node4.

### Steps

1. Disconnect all the network cables from node2.
2. Remove node2 from the AFF A250 chassis.
3. Insert the NVMe shelf module (NSM) into the bay of node2.
4. Connect the NSM to node4 by cabling the node4 100GbE port e0c to NSM B port e0a.

## AFF A400 HA pair with one NS224 shelf



5. Connect the 25GbE cabling from node2 ports e0c and e0d to any two 25GbE onboard ports (e0e, e0f, e0g, or e0h) on node4 to create temporary cluster connections.



If the AFF A400 system uses FC ports as onboard ports, install a 25Gb Ethernet adapter into each node for cluster connectivity during the migration.

6. Connect the 25GbE HA interconnect cables between the AFF A400 nodes using ports e0a and e0b. Do not cross-connect the ports.
7. Connect the 100GbE cluster interconnect cables between the AFF A400 nodes using ports e3a and e3b. Do not cross-connect the ports.

### What's next?

[Reassign drives from node2 to node4](#)

## Reassign drives from node2 to node4

After converting AFF A250 node2 to an NS224 drive shelf and connecting to AFF A400 node4, you must reassign the drives that previously belonged to node2 to node4.

### Before you begin

Verify that both node3 and node4 are on standby at the LOADER prompt.

### About this task

You perform the following steps on node4.

### Steps

1. At the LOADER prompt, boot node4 into Maintenance Mode:

```
boot_ontap maint
```

2. Show the state of the 100GbE interfaces:

```
storage port show
```

3. Set the 100GbE interfaces to storage ports:

```
storage port modify -p e0c -m storage
```

```
storage port modify -p e0d -m storage
```

4. Verify the mode changes to the 100GbE interfaces:

```
storage port show
```

Output like the following example should display:

```
*> storage port modify -p e0c -m storage
Nov 10 16:27:23 [localhost:nvmeof.port.modify:notice]: Changing NVMe-oF
port e0c to storage mode.

Nov 10 16:27:29 [localhost:nvmeof.subsystem.add:notice]: NVMe-oF
subsystem added at address fe80::2a0:98ff:fefa:8885.

*> storage port modify -p e0d -m storage
Nov 10 16:27:34 [localhost:nvmeof.port.modify:notice]: Changing NVMe-oF
port e0d to storage mode.

Nov 10 16:27:38 [localhost:nvmeof.subsystem.add:notice]: NVMe-oF
subsystem added at address fe80::2a0:98ff:fefa:8886.

*> storage port show
Port Type Mode      Speed(Gb/s) State      Status  VLAN ID
---- -
e0c  ENET storage 100 Gb/s   enabled  online  30
e0d  ENET storage 100 Gb/s   enabled  offline 30
```

5. Display all attached drives:

```
disk show -v
```

6. Record the local system ID value; this is the system ID of node4. Also record the system IDs of node1 and node2 from the "OWNER" column.
7. Reassign all drives from node2 to node4:

```
disk reassign -s node2_system_ID -d node4_system_ID -p node1_system_ID
```

8. Verify that all reassigned drives are viewable to the new system ID:

```
disk show -s node4_System_ID
```



If drives are not viewable, **STOP** and contact technical support for assistance.

9. Verify that the root aggregate of node2 is reported in the output and the aggregate is online:

```
aggr status
```

10. Exit maintenance mode:

```
halt
```

## What's next?

[Migrate data aggregates, epsilon, and LIFs on node1 to node4](#)

## Migrate data aggregates, epsilon, and LIFs on node1 to node4

Before converting AFF A250 node1 to a drive shelf, you migrate the data aggregates, epsilon, and logical interfaces (LIFs) on node1 to AFF A400 node4.

### Steps

1. At the LOADER prompt for node4, boot the node into the boot menu:

```
boot_ontap menu
```

2. Select option 6 `Update flash from backup config` to restore the `/var` file system to node4.

This replaces all flash-based configuration with the last backup to disks.

3. Enter `y` to continue.



The node automatically reboots to load the new copy of the `/var` file system.

The node reports a system ID mismatch warning. Enter `y` to override the system ID.

4. Migrate the cluster LIFs:

```
set -privilege advanced
```

```
network port show
```



If the system cluster ports are not similar when upgrading an AFF A250 to an AFF A400, you might have to temporarily change the interfaces on node4 into cluster ports:

```
network port modify -node node4 -port port_name -mtu 9000 -ipSPACE  
Cluster
```

```
network interface migrate -vserver Cluster -lif cluster_LIF -destination-node  
node4 -destination-port port_name
```

5. Wait for the cluster to come into quorum and then verify that the cluster nodes are healthy:

```
- cluster show
```



The HA pair and storage failover remain disabled in the current state.

6. Move the cluster LIFs to the temporary 25G cluster ports on node4:

```
network interface modify
```

7. If interface groups and data VLANs are in use on the AFF A250 cluster you are upgrading, perform this step. If not, go to [Step 8](#).

The physical network port names differ between AFF A250 and AFF A400 systems. As a result, there might be incorrectly configured interface groups and displaced VLANs on node4. Check, and if required, fix any incorrectly configured interface groups and displaced VLANs.

8. Migrate the data aggregates on node1 to node4:

```
storage aggregate relocation start -aggregate-list aggregate_list_name -node  
node1 -destination node4 -ndo-controller-upgrade true -override-destination  
-checks true
```

9. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

10. Migrate the epsilon by removing it from node1 and instead moving it to node4.

a. Remove epsilon from node1:

```
cluster modify -epsilon false -node node1
```

b. Move epsilon to node4:

```
cluster modify -epsilon true -node node4
```

11. Display the cluster status:

```
cluster show
```

12. Display all data network LIFs:

```
network interface show -role data
```

13. Migrate all data LIFs to node4:

```
network interface migrate -vserver vserver_name -lif lif_name -destination  
-node node4 -destination-port port_name
```

14. Display the status of all data LIFs in the cluster:

```
network interface show -role data
```

15. If any LIFs are down, set the administrative status of the LIFs to `up` by entering the following command, once for each LIF:

```
network interface modify -vserver vserver_name -lif lif_name -status-admin up
```

16. Migrate the cluster management LIF:

```
network interface migrate -vserver vserver_name -lif cluster_mgmt -destination  
-node node4 -destination-port port_name
```

17. Display the status of the cluster management LIF:

```
network interface show cluster_mgmt
```

18. Halt node1:

```
halt -node node1 -inhibit-takeover true -ignore-quorum-warnings true
```

### What's next?

[Convert node1 to a drive shelf and connect to node3](#)

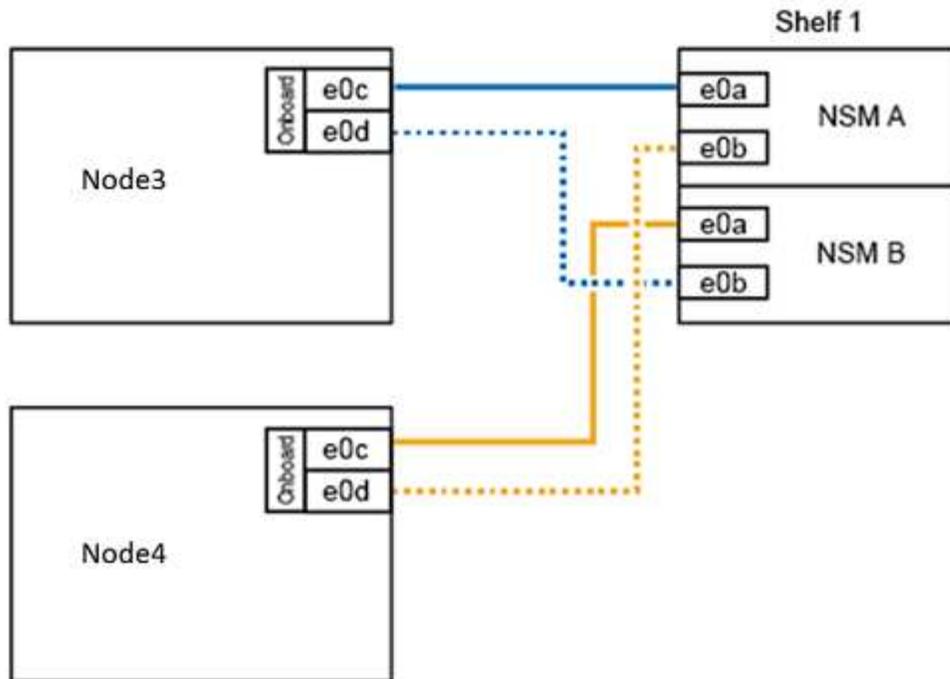
## Convert node1 to a drive shelf and connect to node3

Convert AFF A250 node1 to an NS224 drive shelf and then connect to AFF A400 node3 before reassigning drives from node1 to node3.

### Steps

1. Disconnect all network cables from node1.
2. Remove node1 from the AFF A250 chassis.
3. Insert the NVMe shelf module (NSM) into the bay of node1.
4. Connect the NSM to node3 by cabling node3 100GbE port e0c to NSM A port e0a.

## AFF A400 HA pair with one NS224 shelf



5. Move the temporary cluster connections to node3 by moving the 25GbE cabling from node1 ports e0c and e0d to any two 25GbE onboard ports (e0e, e0f, e0g, or e0h) on node3.



If the AFF A400 system uses FC ports as onboard ports, install a 25Gb Ethernet adapter into each node for cluster connectivity during the migration.

### What's next?

[Reassign drives from node1 to node3](#)

## Reassign drives from node1 to node3

After converting AFF A250 node1 to an NS224 drive shelf and connecting to AFF A400 node3, you must reassign the drives that previously belonged to node1 to node3.

### Steps

1. At the LOADER prompt, boot node3 into Maintenance Mode:

```
boot_ontap maint
```

2. Show the state of the 100GbE interfaces:

```
storage port show
```

3. Set the 100GbE interfaces to storage ports:

```
storage port modify -p e0c -m storage
```

```
storage port modify -p e0d -m storage
```

#### 4. Verify the mode changes to the 100GbE interfaces:

```
storage port show
```

Output like the following example should display:

```
*> storage port modify -p e0c -m storage
Nov 10 16:27:23 [localhost:nvmeof.port.modify:notice]: Changing NVMe-oF
port e0c to storage mode.

Nov 10 16:27:29 [localhost:nvmeof.subsystem.add:notice]: NVMe-oF
subsystem added at address fe80::2a0:98ff:fefa:8885.

*> storage port modify -p e0d -m storage
Nov 10 16:27:34 [localhost:nvmeof.port.modify:notice]: Changing NVMe-oF
port e0d to storage mode.

Nov 10 16:27:38 [localhost:nvmeof.subsystem.add:notice]: NVMe-oF
subsystem added at address fe80::2a0:98ff:fefa:8886.

*> storage port show
Port Type Mode      Speed(Gb/s) State      Status  VLAN ID
----  -
e0c  ENET storage 100 Gb/s   enabled  online  30
e0d  ENET storage 100 Gb/s   enabled  offline 30
```

#### 5. Display all attached drives:

```
disk show -v
```

#### 6. Record the local system ID value; this is the system ID of node3. Also record the system IDs of node1 and node2 from the "OWNER" column.

#### 7. Reassign all drives from node1 to node3:

```
disk reassign -s node1_system_ID -d node3_system_ID -p node4_system_ID
```

#### 8. Verify that all reassigned drives are viewable to the new system ID:

```
disk show -s node3_system_ID
```



If drives are not viewable, **STOP** and contact technical support for assistance.

#### 9. Exit Maintenance Mode:

```
halt
```

### What's next?

[Migrate LIFs and data aggregates on node4 to node3](#)

## Migrate LIFs and data aggregates on node4 to node3

To complete the upgrade, you connect node3 to node4 and then migrate the data logical interfaces (LIFs) and data aggregates on node4 to node3.

### Steps

1. At the LOADER prompt for node3, boot the node into the boot menu:

```
boot_ontap menu
```

2. Select option 6 Update flash from backup config to restore the /var file system to node3.

This replaces all flash-based configuration with the last backup to disks.

3. Enter `y` to continue.

4. Allow the node to boot as normal.



The node automatically reboots to load the new copy of the /var file system.

The node reports a warning that there is a system ID mismatch. Enter `y` to override the system ID.

5. Connect node3 to node4:

- a. Attach multipath high availability (MPHA) cables to the NS224 shelf to ensure redundancy. Connect node3 100GbE port e0d to NSM B port e0b and connect node4 100GbE port e0d to NSM A port e0a.
- b. Verify that HA ports e0a and e0b are connected between the nodes.
- c. Verify that cluster ports e3a and e3b are connected between the nodes.

6. Migrate the cluster LIFs:

```
set -privilege advanced
```

```
network port show
```

7. Modify the cluster broadcast domain to include the desired cluster ports:

```
network port broadcast-domain remove-ports -broadcast-domain  
broadcast_domain_name -ports port_names
```

```
network port broadcast-domain add-ports -broadcast-domain Cluster -ports  
port_names
```



Beginning with ONTAP 9.8, new IPspaces and one or more broadcast domains might be designated to existing physical ports that are intended for cluster connectivity.

8. Modify the cluster IPspace to include the desired cluster ports and set the maximum transmission unit to 9000 if not already set:

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

9. Display all cluster network LIFs:

```
network interface show -role cluster
```

10. Migrate all cluster network LIFs on both nodes to their home ports:

```
network interface migrate -vserver vservice_name -lif lif_name -destination  
-node node_name -destination-port port_name
```

11. Display all cluster network LIFs:

```
network interface show -role cluster
```

12. Verify the home ports for the cluster network LIFs:

```
network interface modify -vserver vservice_name -lif lif_name -home-port  
port_name
```

13. Migrate all data LIFs to node3:

```
network interface migrate -vserver vservice_name -lif lif_name -destination  
-node node_name -destination-port port_name
```

14. Display all data network LIFs:

```
network interface show -role data
```

15. Configure the home node and home port for all data LIFs. If any LIFs are down, set the administrative status of the LIFs to up by entering the following command, once for each LIF:

```
network interface modify -vserver vservice_name -lif lif_name -home-node  
node_name -home-port port_name -status-admin up
```

16. Migrate the cluster management LIF:

```
network interface migrate -vserver vservice_name -lif cluster_mgmt -destination  
-node node3 -destination-port port_name
```

17. Display the status of the cluster management LIF:

```
network interface show cluster_mgmt
```

18. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

19. Enable the HA pair, storage failover, and auto-giveback:

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

20. Migrate data aggregates owned by node4 to node3:

```
storage aggregate relocation start -aggregate aggregate_name -node node4  
-destination node3
```

21. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

22. Enable auto-revert of the network LIFs across the nodes:

```
network interface modify -lif * -auto-revert true
```

23. Enable storage failover automatic giveback:

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

24. Display the cluster status:

```
cluster show
```

25. Display failover eligibility:

```
storage failover show
```



In the cluster report output, a node might incorrectly own aggregates that belong to another node. If this occurs, normalize by performing a takeover and giveback from both sides of the cluster.

26. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

## Upgrade from FAS2820

### Upgrade from FAS2820 by converting to DS212C drive shelf

Perform a nondisruptive upgrade of a NetApp FAS2820 system by converting each node to a DS212C drive shelf. Then, connect the drive shelves to the replacement nodes. This moves the FAS2820 onboard storage to the replacement system.

#### Before you begin

Review the general upgrade scenarios and upgrade considerations for upgrading by moving volumes or storage:

- [Decide whether to upgrade by moving volumes or storage](#)
- [Considerations for upgrading controller hardware](#)

#### About this task

The FAS2820 high-availability (HA) pair controllers are node1 and node2 and the replacement HA pair controllers are node3 and node4.



#### Migrate LIFs and data aggregates on node2 to node1

Before converting FAS2820 node2 to a drive shelf, migrate the logical interfaces (LIFs) and data aggregates on

node2 to node1.

2

### Convert node2 to a drive shelf and connect to node4

Convert FAS2820 node2 to a DS212C drive shelf and then connect to node4 before reassigning drives from node2 to node4.

3

### Reassign drives from node2 to node4

After converting FAS2820 node2 to a DS212C drive shelf and connecting to node4, reassign the drives that previously belonged to node2 to node4.

4

### Migrate data aggregates, epsilon, and LIFs on node1 to node4

Before converting FAS2820 node1 to a drive shelf, migrate the data aggregates, epsilon, and LIFs on node1 to node4.

5

### Convert node1 to a drive shelf and connect to node3

Convert FAS2820 node1 to a DS212C drive shelf and connect to node3 before reassigning drives from node1 to node3.

6

### Reassign drives from node1 to node3

After converting FAS2820 node1 to a DS212C drive shelf and connecting to node3, reassign the drives that previously belonged to node1 to node3.

7

### Migrate LIFs and data aggregates on node4 to node3

To complete the upgrade, connect node3 to node4 and then migrate the data LIFs and data aggregates on node4 to node3.

## Migrate LIFs and data aggregates on FAS2820 node2 to node1

Migrate the logical interfaces (LIFs) and data aggregates on FAS2820 node2 to node1 before converting node2 to a drive shelf.

### Before you begin

Verify that you meet the following requirements:

- The FAS2820 and replacement controllers have the same ONTAP release and patch version, when possible. Refer to the [NetApp Hardware Universe](#) for supported ONTAP releases.



- You need to netboot and install the ONTAP version on node3 and node4 to match the ONTAP version on the FAS2820 systems. Node3 and node4 are the replacement controllers.
- Both the primary and backup boot image of the node3 and node4 controllers must have the same ONTAP version.
- You need to clear any residual cluster configuration on node3 and node4 by performing a `wipeconfig` from the boot menu.

- Both replacement controllers are on standby at the LOADER prompt.
- All required cabling is available.

### About this task

The following steps are performed on FAS2820 node1.

### Steps

1. Access the advanced privilege level:

```
set -privilege advanced
```

2. Disable storage failover automatic giveback:

```
storage failover modify -node node1 -auto-giveback false
```

3. Disable auto-revert of the LIFs across both nodes of the HA pair:

```
network interface modify -lif * -auto-revert false
```

4. Display the status of all data network LIFs:

```
network interface show -role data
```

5. Display the status of the cluster management LIFs:

```
network interface show -role cluster-mgmt
```

6. Migrate all data LIFs from the storage virtual machines hosted on node2:

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



This command only migrates non-SAN LIFs. You cannot use it to migrate iSCSI and FCP LIFs.

7. Display the status of all data LIFs in the cluster:

```
network interface show -role data
```

8. If any LIFs are down, set their administrative status to up by entering the following command once for each LIF:

```
network interface modify -vserver <vserver_name> -lif <lif_name> -status  
-admin up
```

9. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

10. Display failover eligibility:

```
storage failover show
```

11. Migrate the data aggregates on node2 to node1:

```
storage aggregate relocation start -aggregate <aggregate_name> -node  
<node2> -destination <node1>
```

12. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

13. Display the status of all data volumes in the cluster:

```
volume show
```

14. Display the ha status and ownership of epsilon:

```
cluster show
```

15. Disable cluster ha:

```
cluster ha modify -configured false
```

16. Display the ha status and ownership of epsilon:

```
cluster show
```

17. Halt node2:

```
halt -node <node2> -inhibit-takeover true -ignore-quorum-warnings true
```

### What's next?

[Convert node2 to a drive shelf and connect to node4](#)

## Convert FAS2820 node2 to a drive shelf and connect to node4

Convert FAS2820 node2 to a DS212C drive shelf. Connect to node4 before reassigning drives from node2 to node4. Node1 and node2 are the controllers inside the DS212C shelf.

### Steps

1. Disconnect all the network cables from node2.
2. Remove node2 from the FAS2820 chassis.
3. Insert the IOM12 or IOM12B module into the bay of node2.
4. Cable the node4 serial-attached SCSI (SAS) port to an available port on the IOM12 or IOM12B module. See the [Hardware Universe](#) to verify the SAS ports for your system.
5. Create temporary cluster connections by connecting node1 ports e0a and e0b to any two 25GbE ports on node4.



If node4 only supports 10GbE cluster connections, you must have 10GbE cabling to create temporary cluster connections.

### What's next?

[Reassign drives from node2 to node4](#)

## Reassign drives from FAS2820 node2 to node4

Reassign the drives that previously belonged to FAS2820 node2 to node4.

### Before you begin

Verify that both node3 and node4 are on standby at the LOADER prompt.

### About this task

Perform the following steps on node4.

## Steps

1. At the LOADER prompt, boot node4 into Maintenance Mode:

```
boot_ontap maint
```

The Maintenance Mode prompt appears.

2. Display all attached drives:

```
disk show -v
```

3. Record the local system ID value; this is the system ID of node4. Also record the system IDs of node1 and node2 from the "OWNER" column.
4. Reassign all drives from node2 to node4.

### If you're using whole disks

Run the following command:

```
disk reassign -s <node2_system_ID> -d <node4_system_ID>
```

### If you're using partitioned disks

Run the following command:

```
disk reassign -s <node2_system_ID> -d <node4_system_ID> -p  
<node1_system_ID>
```

5. Verify that all reassigned drives are viewable to the new system ID:

```
disk show -s <node4_System_ID>
```



If drives are not viewable, **stop** and contact technical support for assistance.

6. Verify that the root aggregate of node2 is reported in the output and the aggregate is online:

```
aggr status
```

7. Exit maintenance mode:

```
halt
```

### What's next?

[Migrate data aggregates, epsilon, and LIFs on node1 to node4](#)

## Migrate data aggregates, epsilon, and LIFs on FAS2820 node1 to node4

Migrate the data aggregates, epsilon, and logical interfaces (LIFs) on the FAS2820 node1 to node4.

### About this task

Perform the following steps on node4.

### Steps

1. At the LOADER prompt for node4, set partner-system ID:

```
setenv partner-sysid <system_ID_of_node1>
```

2. Check the partner system-ID:

```
printenv partner-sysid
```

3. Save the changes:

```
saveenv
```

4. Boot the node into the boot menu:

```
boot_ontap menu
```

5. At the boot menu, select option 6 Update flash from backup config to restore the /var file system to node4.

This replaces all flash-based configuration with the last backup to disks.

6. Enter `y` to continue.



The node automatically reboots to load the new copy of the /var file system.

The node reports a system ID mismatch warning. Enter `y` to override the system ID.

7. Migrate the cluster LIFs:

```
set -privilege advanced
```

```
network port show
```

If the system cluster ports are not similar when upgrading a FAS2820 to a replacement controller, you might have to temporarily change the interfaces on node4 into cluster ports:



```
network port modify -node <node4> -port <port_name> -mtu 9000  
-ipSPACE Cluster
```

```
network interface migrate -vserver Cluster -lif <cluster_LIF>  
-destination-node <node4> -destination-port <port_name>
```

8. Wait for the cluster to reach quorum, then verify that the cluster nodes are healthy:

```
- cluster show
```



The HA pair and storage failover remain disabled in the current state.

9. Move the cluster LIFs to the temporary 25G cluster ports on node4:

```
network interface modify -vserver Cluster -lif <cluster_LIF> -home-node  
<node4> -home-port <port_name>
```

10. Complete this step only if interface groups and data VLANs are used on the FAS2820 cluster you are upgrading. Otherwise, go to [Step 11](#).

The physical network port names on the replacement controller differ from those on the FAS2820. This might lead to displaced VLANs and incorrectly configured interface groups on node4.

a. Show the displaced VLANs:

```
cluster controller-replacement network displaced-vlans show
```

b. Restore the displaced VLANs:

```
cluster controller-replacement network displaced-vlans restore
```

- c. Fix the incorrectly configured interface groups. The port names between the FAS2820 and the controllers you are upgrading might be different. Update the interface groups with the correct member ports:

```
ifgrp remove-port -node <node2> -ifgrp <ifgrp_name> -port <port_name>
```

```
ifgrp add-port -node <node2> -ifgrp <ifgrp_name> -port <port_name>
```

11. Migrate the data aggregates on node1 to node4:

```
storage aggregate relocation start -aggregate-list <aggregate_list_name>  
-node <node1> -destination <node4> -ndo-controller-upgrade true  
-override-destination-checks true
```

12. Show data aggregate status:

```
storage aggregate show
```

13. Migrate the epsilon by removing it from node1 and moving it to node4.

- a. Remove epsilon from node1:

```
cluster modify -epsilon false -node <node1>
```

- b. Move epsilon to node4:

```
cluster modify -epsilon true -node <node4>
```

14. Show cluster status to verify the epsilon change:

```
cluster show
```

15. Display all data network LIFs:

```
network interface show -role data
```

16. Migrate all data LIFs to node4:

```
network interface migrate -vserver <vserver_name> -lif <lif_name>
-destination-node <node4> -destination-port <port_name>
```

17. Display the status of all data LIFs in the cluster:

```
network interface show -role data
```

18. If any LIFs are down, set their administrative status to up by entering the following command once for each LIF:

```
network interface modify -vserver <vserver_name> -lif <lif_name> -status
-admin up
```

19. Migrate the cluster management LIF:

```
network interface migrate -vserver <vserver_name> -lif
<cluster_mgmt_lif> -destination-node <node4> -destination-port
<port_name>
```

20. Display the status of the cluster management LIF:

```
network interface show -role cluster-mgmt
```

21. Halt node1:

```
halt -node <node1> -inhibit-takeover true -ignore-quorum-warnings true
```

## What's next?

[Convert node1 to a drive shelf and connect to node3](#)

## Convert FAS2820 node1 to a drive shelf and connect to node3

Convert FAS2820 node1 to a DS212C drive shelf. Connect to node3 before reassigning drives from node1 to node3..

### Steps

1. Disconnect all network cables from node1.
2. Remove node1 from the FAS2820 chassis.
3. Insert the IOM12 or IOM12B module into the bay of node1.
4. Cable the node3 SAS port to an available port on the IOM12 or IOM12B module. See the [Hardware](#)

[Universe](#) to verify the SAS ports for your system. See the [Hardware Universe](#) to verify the SAS ports for your system.

5. Create temporary cluster connections by connecting the node4 cluster ports to any cluster ports on node3.



If node3 only supports 10GbE cluster connections, you must have 10GbE cabling to create temporary cluster connections.

### What's next?

[Reassign drives from node1 to node3](#)

## Reassign drives from FAS2820 node1 to node3

Reassign the drives that were previously assigned to FAS2820 node1 to node3.

### About this task

Perform the following steps on node3.

### Steps

1. At the LOADER prompt, boot node3 into Maintenance Mode:

```
boot_ontap maint
```

The Maintenance Mode prompt appears.

2. Display all attached drives:

```
disk show -v
```

3. Record the local system ID value; this is the system ID of node3. Also record the system IDs of node1 and node4 from the "OWNER" column.
4. Reassign all drives from node1 to node3:

#### If you're using whole disks

Run the following command:

```
disk reassign -s <node1_system_ID> -d <node3_system_ID>
```

#### If you're using partitioned disks

Run the following command:

```
disk reassign -s <node1_system_ID> -d <node3_system_ID> -p  
<node4_system_ID>
```

5. Verify that all reassigned drives are viewable to the new system ID:

```
disk show -s <node3_system_ID>
```



If drives are not viewable, **stop** and contact technical support for assistance.

6. Exit Maintenance Mode:

```
halt
```

### What's next?

[Migrate LIFs and data aggregates on node4 to node3](#)

## Migrate LIFs and data aggregates on FAS2820 node4 to node3

To complete the upgrade, you connect FAS2820 node3 to node4 and then migrate the data logical interfaces (LIFs) and data aggregates on node4 to node3.

### About this task

Perform the following steps on node3.

### Steps

1. At the LOADER prompt for node3, boot the node into the boot menu:

```
boot_ontap menu
```

2. Select option 6 Update flash from backup config to restore the /var file system to node3.

This replaces all flash-based configuration with the last backup to disks.

3. Enter *y* to continue.

4. Allow the node to boot as normal.



The node automatically reboots to load the new copy of the /var file system.

The node reports a warning that there is a system ID mismatch. Enter *y* to override the system ID.

5. Verify that the cluster and HA ports are connected between node3 to node4.

6. Display the cluster and HA ports on node3 and node4:

```
set -privilege advanced
```

```
network port show
```

7. Modify the cluster broadcast domain to include the desired cluster ports:

```
network port broadcast-domain remove-ports -broadcast-domain  
<broadcast_domain_name> -ports <port_names>
```

```
network port broadcast-domain add-ports -broadcast-domain Cluster -ports  
<port_names>
```



Beginning with ONTAP 9.8, new IPspaces and one or more broadcast domains might be designated to existing physical ports that are intended for cluster connectivity.

8. Modify the cluster IPspace to include the desired cluster ports and set the maximum transmission unit to 9000 if not already set:

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

9. Display all cluster network LIFs:

```
network interface show -role cluster
```

10. Migrate all cluster network LIFs on both nodes to their planned home ports:

```
network interface migrate -vserver <vserver_name> -lif <lif_name>  
-destination-node <node_name> -destination-port <port_name>
```

11. Display all cluster network LIFs:

```
network interface show -role cluster
```

12. Configure the home ports for the cluster network LIFs:

```
network interface modify -vserver <vserver_name> -lif <lif_name> -home  
-port <port_name>
```

13. Migrate all data LIFs meant for node3 back to node3:

```
network interface migrate -vserver <vserver_name> -lif <lif_name>
-destination-node <node3> -destination-port <port_name>
```

14. Display all data network LIFs:

```
network interface show -role data
```

15. Configure the home node and home port for all data LIFs. If any LIFs are down, set their administrative status to up by entering the following command once for each LIF:

```
network interface modify -vserver <vserver_name> -lif <lif_name> -home
-node <node_name> -home-port <port_name> -status-admin up
```

16. Migrate the cluster management LIF:

```
network interface migrate -vserver <vserver_name> -lif
<cluster_mgmt_lif> -destination-node <node3> -destination-port
<port_name>
```

17. Display the status of the cluster management LIF:

```
network interface show -role cluster-mgmt
```

18. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

19. Enable cluster high availability in the two-node cluster:

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

20. Enable and verify storage failover for node3 and node4:

```
storage failover modify -node <node3> -enabled true
```

```
storage failover modify -node <node4> -enabled true
```

```
storage failover show
```

21. Migrate data aggregates owned by node4 that should be owned by node3:

```
storage aggregate relocation start -aggregate <aggregate_name> -node  
<node4> -destination <node3>
```

22. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

23. Enable auto-revert of the network LIFs across the nodes:

```
network interface modify -lif * -auto-revert true
```

24. Enable storage failover automatic giveback:

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

25. Display the cluster status:

```
cluster show
```

26. Display failover eligibility:

```
storage failover show
```



In the cluster report output, a node might incorrectly own aggregates that belong to another node. If this occurs, perform a takeover and giveback from both sides of the cluster.

27. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

## Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

## Trademark information

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