



---

# Using aggregate relocation to upgrade controller hardware on a pair of nodes running ONTAP 9.x

You can nondisruptively upgrade the controller hardware on a pair of nodes running ONTAP 9.x by migrating non-root aggregates from the original nodes to new nodes in the same cluster. The data hosted on the nodes that are being upgraded is accessible during the upgrade.

## About this task

This procedure uses a method called aggregate relocation (ARL), which takes advantage of the HA configuration to enable you to move ownership of non-root aggregates from one node to another, if they share storage within the same cluster.

During the procedure, you upgrade the original controller hardware with the replacement controller hardware, relocating the ownership of non-root aggregates. You migrate aggregates multiple times from node to node to ensure that at least one node is serving data from the aggregates throughout the upgrade procedure. You also migrate data logical interfaces (LIFs) and assign the network ports on the new controller to the interface groups as you proceed.

**Important:** This procedure is complex and assumes that you have advanced ONTAP administration skills. You also should read and understand the “How the upgrade using ARL works” and “Guidelines for upgrading controllers with ARL” sections before beginning the upgrade.

In this document, 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*.

**Attention:** This procedure assumes that the replacement controller hardware is new and has not been used. The steps required to prepare used controllers with the `wipeconfig` command are not included in this procedure. You must contact technical support if the replacement controller hardware was previously used, especially if the controllers were running Data ONTAP in 7-Mode.

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

**Note:** You can use this procedure to upgrade the controller hardware in clusters with more than two nodes; however, you need to perform the procedure separately for each HA pair in the cluster.

This procedure applies both to FAS systems, V-Series systems, All Flash FAS systems, and systems with FlexArray Virtualization Software.

**Note:** The new FAS systems released after ONTAP 8.2.1 can attach to storage arrays if the required license is installed. No new *V-Series system* are released starting with Data ONTAP 8.2.1. The existing V-Series systems are supported in Data ONTAP 8.2.1. See the V-Series Support Matrix at *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com) for information about the storage array and V-Series models.

This document uses the term *systems with FlexArray Virtualization Software* to refer to systems that belong to these new platforms. It uses the term *V-Series system* to refer to the separate hardware systems that can attach to storage arrays.

## Deciding whether to use the aggregate relocation guide

This guide describes how to upgrade the storage controllers in an HA pair with new controllers while keeping all the existing data and disks. This is a complex procedure that should be used only by experienced administrators.

You should use this guide under the following circumstances:

- You do not want to add the new controllers as a new HA pair to the cluster and migrate the data using volume moves.
- You are experienced in administering ONTAP and are comfortable with the risks of working in the diagnostic privilege mode.

If you prefer a different method of upgrading the controller hardware and are willing to do volume moves, use the *Controller Hardware Upgrade Express Guide* instead. See the ONTAP 9 Documentation Center to access ONTAP 9 product documentation.

### Related information

[Controller Hardware Upgrade Express Guide](#)  
[ONTAP 9 Documentation Center](#)

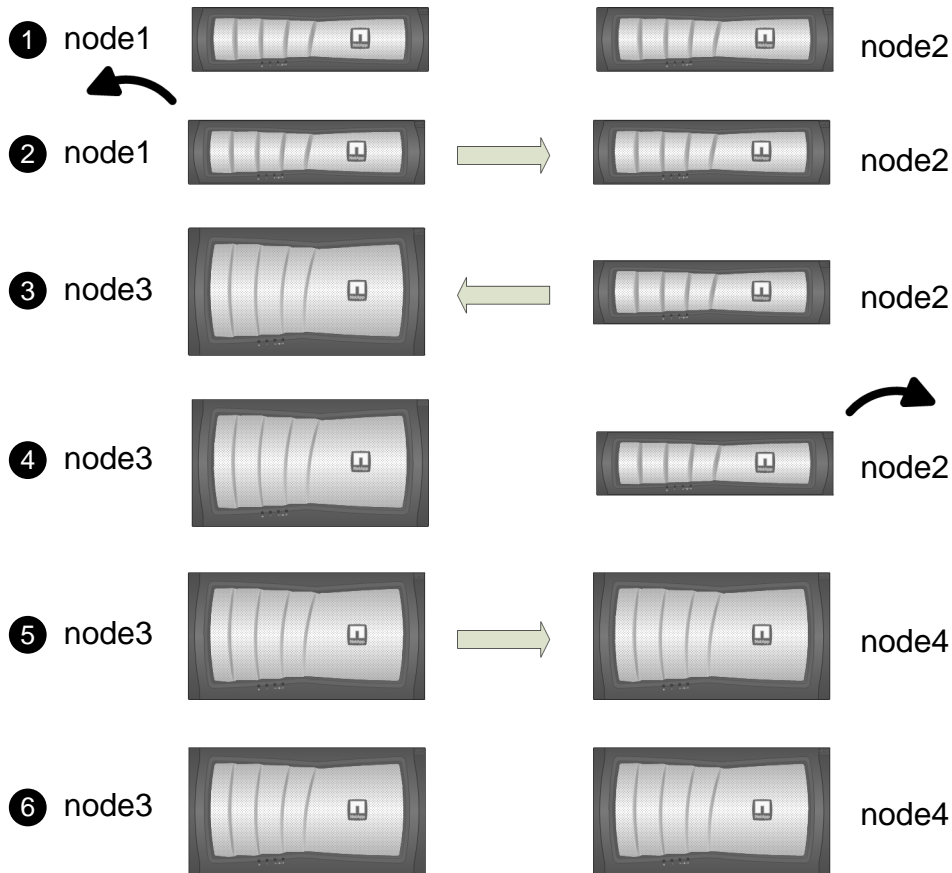
## How the upgrade using ARL works

Before you upgrade the nodes using ARL, you should understand how the procedure works. In this document, the procedure is broken down into several stages.

During the procedure, you upgrade the original controller hardware with the replacement controller hardware, one controller at a time, taking advantage of the HA configuration to relocate the ownership of non-root aggregates. All non-root aggregates must undergo two relocations to reach their final destination, which is the correct upgraded node.

Each aggregate has a home owner and current owner. The home owner is the actual owner of the aggregate, and the current owner is the temporary owner.

The following illustration shows the stages of the procedure. The thick, light gray arrows represent the relocation of aggregates and the movement of LIFs, and the thinner black arrows represent the removal of the original nodes. The smaller controller images represent the original nodes, and the larger controller images represent the new nodes.



The following table describes the high-level tasks you perform during each stage and the state of aggregate ownership at the end of the stage. Detailed steps are provided later in the procedure:

<b>1</b>	<p>Stage 1: Preparing for the upgrade</p> <ol style="list-style-type: none"> <li>1. Determining whether the controller has aggregates on internal disk drives. This step is required only if you are upgrading from a controller with internal disk drive.</li> <li>2. Preparing the nodes for the upgrade</li> <li>3. Rekeying disks for Storage Encryption This task is required only if you are upgrading from a system with self-encrypting drives.</li> <li>4. Verifying the SnapMirror relationship state on the cluster and quiesce all relationships between the clusters.</li> <li>5. Preparing for netboot</li> </ol> <p>Aggregate ownership at the end of Stage 1:</p> <ul style="list-style-type: none"> <li>• Node1 is the home owner and current owner of node1 aggregates.</li> <li>• Node2 is the home owner and current owner of node2 aggregates.</li> </ul>
----------	--

<p style="text-align: center;"><b>2</b></p>	<p>Stage 2: Retiring node1</p> <ol style="list-style-type: none"> <li>1. Relocating non-root aggregates from node1 to node2</li> <li>2. Moving non-SAN data LIFs owned by node1 to node2</li> <li>3. Recording node1 information</li> <li>4. Retiring node1</li> </ol> <p>Aggregate ownership at the end of Stage 2:</p> <ul style="list-style-type: none"> <li>• Node1 is the home owner of node1 aggregates.</li> <li>• Node2 is the current owner of node1 aggregates.</li> <li>• Node2 is the home owner and current owner of node2 aggregates.</li> </ul>
<p style="text-align: center;"><b>3</b></p>	<p>Stage 3: Bringing up node3</p> <ol style="list-style-type: none"> <li>1. Installing and booting node3</li> <li>2. Setting the FC or CNA configuration on node3</li> <li>3. Mapping ports from node1 to node3</li> <li>4. Verifying the node3 installation</li> <li>5. Moving non-SAN data LIFs owned by node1 from node2 to node3 and verifying SAN LIFs on node3</li> <li>6. Relocating non-root aggregates from node2 to node3</li> <li>7. Moving non-SAN data LIFs owned by node2 to node3</li> </ol> <p>Aggregate ownership at the end of Stage 3:</p> <ul style="list-style-type: none"> <li>• Node2 is the home owner of node2 aggregates but not the current owner.</li> <li>• Node3 is the home owner and current owner of aggregates that originally belonged to node1.</li> <li>• Node3 is the current owner of aggregates belonging to node2 but not the home-owner.</li> </ul>
<p style="text-align: center;"><b>4</b></p>	<p>Stage 4: Retiring node2</p> <ol style="list-style-type: none"> <li>1. Recording node2 information</li> <li>2. Retiring node2</li> </ol> <p>No changes occur in aggregate ownership.</p>

<b>5</b>	<p>Stage 5: Bringing up node4</p> <ol style="list-style-type: none"> <li>1. Installing and booting node4</li> <li>2. Setting the FC or CNA configuration on node4</li> <li>3. Mapping ports from node2 to node4</li> <li>4. Verifying node4 installation</li> <li>5. Moving non-SAN data LIFs owned by node2 from node3 to node4 and verifying SAN LIFs on node4</li> <li>6. Relocating node2's non-root aggregates from node3 to node4</li> </ol> <p>Aggregate ownership at the end of Stage 5:</p> <ul style="list-style-type: none"> <li>• Node3 is the home owner and current owner of aggregates that originally belonged to node1.</li> <li>• Node4 is the home owner and current owner of aggregates that originally belonged to node2.</li> </ul>
<b>6</b>	<p>Stage 6: Completing the upgrade</p> <ol style="list-style-type: none"> <li>1. Ensuring that the new controllers are set up correctly</li> <li>2. Setting up Storage Encryption on the new nodes This task is required only if you are upgrading to a system with self-encrypting drives.</li> <li>3. Decommissioning the old system</li> <li>4. Resume SnapMirror relationship</li> </ol> <p><b>Note:</b> Storage Virtual Machine (SVM) disaster recovery updates will not be interrupted as per the schedules assigned.</p> <p>No changes occur in aggregate ownership.</p>

## Guidelines for upgrading controllers with ARL

To understand whether you can use aggregate relocation (ARL) to upgrade a pair of controllers running ONTAP 9.x depends on the platform and the configuration of both the original and replacement controllers.

### Supported upgrades for ARL

You can upgrade a pair of nodes using ARL under the following circumstances:

- Both the original controllers and the replacement controllers must be running the same version of ONTAP 9.x before the upgrade.
- The replacement controllers must have equal or higher capacity than the original controllers.  
Equal or higher capacity refers to attributes, such as the NVRAM size, volume, LUN, or aggregate count limits; it also refers to the maximum volume or aggregate sizes of the new nodes.
- You can upgrade the following type of systems:
  - A FAS system to a FAS system
  - A FAS system to a V-Series system or a system with FlexArray Virtualization Software
  - A V-Series system to a V-Series system or a system with FlexArray Virtualization Software

- A V-Series system or a system with FlexArray Virtualization Software to a FAS system, provided that the V-Series system or system with FlexArray Virtualization Software has no array LUNs.
- An AFF system to an AFF system

You should verify whether the ARL can be performed on the original and replacement controllers. You should check the size of all defined aggregates and number of disks supported by the original system. Then compare them with the aggregate size and number of disks supported by the new system. This information is available in the *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com). The aggregate size and the number of disks supported by the new system must be equal to or greater than the aggregate size and number of disks supported by the original system. You should validate in the cluster mixing rules whether new nodes can become part of the cluster with the existing nodes, when the original controller is replaced.

**Note:** If the controllers are running ONTAP 8.3 or later and one or both are All Flash FAS models, you should verify that both controllers have the same All Flash Optimized personality set (`system node show -instance node_name`). Both nodes must either have the personality enabled or disabled; you cannot combine a node with the All Flash Optimized personality enabled with a node that does not have the personality enabled in the same HA pair. If the personalities are different, see [KB Article 1015157](#) in the NetApp Knowledge Base for instructions about how to synchronize node personality.

A system can be in a single-chassis HA pair or a dual-chassis HA pair. A single-chassis HA pair is one in which both controller modules are in the same chassis; a dual-chassis HA pair is one in which each controller module is in a separate chassis.

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

## Upgrades not supported for ARL

You cannot perform the following upgrades:

- To or from controllers that cannot run ONTAP 9.x.  
Controller upgrades for systems running Data ONTAP operating in 7-Mode are available on the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com).
- To replacement controllers that do not support the disk shelves connected to the original controllers  
See the *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com) for disk-support information.
- From controllers with root aggregates or data aggregates on internal drives  
If you want to upgrade controllers with root aggregates or data aggregates on internal disk drives, see the procedure *Upgrading a pair of nodes running clustered Data ONTAP by moving volumes* in *Controller Hardware Upgrade Express Guide*

**Note:** If you want to upgrade ONTAP on nodes in a cluster, see the *ONTAP 9 Upgrade and Revert/Downgrade Guide* on the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com).

## Assumptions and terminology

This document is written with the following assumptions:

- You read and understand the guidelines for upgrading the pair of nodes.
- The replacement controller hardware is new and has not been used.

**Note:** If either of the new controllers was used previously, you should disconnect the NVRAM battery.

**Attention:** Do not try to clear the NVRAM contents. If there is a need to clear the contents of NVRAM, contact NetApp technical support.

- You are performing the appropriate command before and after the `modify` commands and comparing the output of both `show` commands to verify that the `modify` command was successful.
- If you have a SAN configuration, you have local and remote LIFs for each Storage Virtual Machine (SVM), on the HA pair.

If you do not have local and remote LIFs for each SVM, you should add the SAN data LIF on the remote and local node for that SVM before beginning the upgrade.

- If you have port sets in a SAN configuration, you must have verified that each bound port set contains at least one LIF from each node in the HA pair.

This procedure uses the term *boot environment prompt* to refer to the prompt on a node from which you can perform certain tasks, such as rebooting the node and printing or setting environmental variables. The prompt is sometimes referred to informally as the *boot loader prompt*.

The boot environment prompt is shown in the following example:

```
LOADER>
```

## Licensing in ONTAP 9.x

When you set up a cluster, the setup wizard prompts you to enter the cluster-base license key. However, some features require additional licenses, which are issued as *packages* that include one or more features. Each node in the cluster must have its own key for each feature to be used in the cluster.

If you do not have new license keys, currently licensed features in the cluster are available to the new controller. However, using unlicensed features on the controller might put you out of compliance with your license agreement, so you should install the new license key or keys for the new controller after the upgrade is complete.

Starting with Data ONTAP 8.2.x, all license keys are 28 uppercase alphabetic characters in length. You can obtain new 28-character license keys for ONTAP 9.x on the NetApp Support Site at [mysupport.netapp.com](https://mysupport.netapp.com). The keys are available in the *My Support* section under *Software licenses*. If the site does not have the license keys you need, you can contact your NetApp sales representative.

For detailed information about licensing, see the *System Administration Reference* and the KB article *Data ONTAP 8.2 and 8.3 Licensing Overview and References* on the NetApp Support Site.

## Storage Encryption

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

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

## Two-node switchless clusters

If you are upgrading nodes in a two-node switchless cluster, you can leave the nodes in the switchless cluster while performing the upgrade. You do not need to convert them to a switched cluster.

## Troubleshooting

If any problems occur while upgrading the controllers, you can refer to the *Troubleshooting* section at the end of the procedure for more information and possible solutions.

If you do not find a solution to the problem you encountered, you should contact technical support.

## Related information

[Controller Hardware Upgrade Express Guide](#)

## Required tools and documentation

You must have specific tools to install the new hardware, and you need to reference other documents during the upgrade process. You also need to record information essential to completing the controller upgrade; a worksheet is provided to record information.

You need the following tools to perform the upgrade:

- Grounding strap
- #2 Phillips screwdriver

You need information in the following documents, which are available from the NetApp Support Site at [mysupport.netapp.com](https://mysupport.netapp.com)

**Note:** If you cannot access the NetApp Support Site when you are installing the new node, download the appropriate documents before beginning this procedure.

Document	Contents
<i>ONTAP 9 Software Setup Guide</i>	Describes how to set up and configure NetApp systems
<i>ONTAP 9 System Administration Reference</i>	Contains instructions for obtaining information to set up the Service Processors (SPs) and performing other administrative tasks
<i>ONTAP 9 Upgrade and Revert/Downgrade Guide</i>	Contains instructions for downloading and upgrading ONTAP
<i>ONTAP 9 Disks and Aggregates Power Guide</i>	Describes how to manage physical storage resources, using disks, RAID groups, and aggregates, and provides detailed information about ARL and Storage Encryption
<i>ONTAP 9 Logical Storage Management Guide</i>	Describes how to efficiently manage your logical storage resources, using volumes, FlexClone volumes, files and LUNs, FlexCache volumes, deduplication, compression, qtrees, and quotas
<i>Installation and Setup Instructions for the model of the new nodes</i>	Contains instructions for installing and cabling the new system.
<i>FlexArray Virtualization Installation Requirements and Reference Guide</i>	Contains cabling instructions and other information for FlexArray Virtualization systems
<i>Hardware Universe at <a href="https://hwu.netapp.com">hwu.netapp.com</a></i>	Contains information about system LEDs and system messages to aid troubleshooting
The appropriate disk shelf guide	Contains instructions for installing and monitoring disk shelves and replacing disk shelf devices
<i>ONTAP 9 SAN Administration Guide</i>	Describes how to configure and manage iSCSI and FC protocols for SAN environments
<i>ONTAP 9 SAN Configuration Guide</i>	Contains information about FC and iSCSI topologies and wiring schemes
<i>ONTAP 9 High-Availability Configuration Guide</i>	Contains cabling instructions and other information for HA pairs
<i>MetroCluster Installation and Configuration Guide</i>	Describes how to install and configure the MetroCluster hardware and software components.



The NetApp Support Site also contains documentation about NICs and other hardware that you might use with your system. It also contains the *Hardware Universe*, which provides information about the hardware that the new system supports. See the ONTAP 9 Documentation Center to access ONTAP 9 product documentation.

**Related information**

[ONTAP 9 Documentation Center](#)

**Worksheet: Information to collect before and during the controller upgrade**

You must gather certain information to successfully upgrade the original nodes. The information includes node IDs, port and LIF details, licensing keys, and IP addresses.

You can use the following worksheet to record information for use later in the procedure:

Information needed	When collected	When used	Collected information
Model, system ID, serial number of original nodes	Stage 1: <i>Preparing the nodes for the upgrade</i>	Stage 3: <i>Installing and booting node3</i> Stage 5: <i>Installing and booting node4</i> Stage 6: <i>Decommissioning the old system</i>	
Shelf and disk information, flash storage details, memory, NVRAM, and adapter cards on original nodes	Stage 1: <i>Preparing the nodes for the upgrade</i>	Throughout the procedure	
Online aggregates and volumes on original nodes	Stage 1: <i>Preparing the nodes for the upgrade</i>	Throughout the procedure to verify that aggregates and volumes remain online except during brief relocation	
Output of commands <code>network port vlan show</code> and <code>network port ifgrp show</code>	Stage 1: <i>Preparing the nodes for the upgrade</i>	Stage 3: <i>Mapping ports from node1 to node3</i> Stage 5: <i>Mapping ports from node2 to node4</i>	
(SAN environments only) Default configuration of FC ports	Stage 1: <i>Preparing the nodes for the upgrade</i>	When configuring FC ports on the new nodes	
(V-Series systems or systems with FlexArray Virtualization software only) Topology for V-Series systems or systems with FlexArray Virtualization software	Stage 1: <i>Preparing the nodes for the upgrade</i>	Stage 3: <i>Installing and booting node3</i> Stage 5: <i>Installing and booting node4</i>	
IP address of SPs	Stage 1: <i>Preparing the nodes for the upgrade</i>	Stage 6: <i>Ensuring that the new controllers are set up correctly</i>	

<b>Information needed</b>	<b>When collected</b>	<b>When used</b>	<b>Collected information</b>
License keys	Stage 1: <i>Preparing the nodes for the upgrade</i>	Stage 6: <i>Ensuring that the new controllers are set up correctly</i>	
IP address for the external key management server	Stage 1: <i>Rekeying disks for Storage Encryption</i>	Stage 6: <i>Setting up Storage Encryption on the new nodes</i>	
Name and path of web-accessible directory where you download files to netboot the nodes	Stage 1: <i>Preparing to netboot</i>	Stage 3: <i>Installing and booting node3</i> Stage 5: <i>Installing and booting node4</i>	
Non-SAN data LIFs owned by node1	Stage 2: <i>Moving non-SAN data LIFs owned by node1 to node2</i>	Later in the section	
Cluster, intercluster, node-management, cluster-management, and physical ports	Stage 2: <i>Recording node1 information</i>	Stage 3: <i>Installing and booting node3</i> Stage 3: <i>Mapping ports from node1 to node3</i>	
Ports on new nodes	Stage 3: <i>Mapping ports from node1 to node3</i>	Later in the section and in the section <i>Mapping ports from node2 to node4</i>	
Available ports and broadcast domains on node3	Stage 3: <i>Mapping ports from node1 to node3</i>	Later in the section	
Non-SAN data LIFs not owned by node2	<i>Moving non-SAN data LIFs belonging to node1 from node2 to node3 and verifying SAN LIFs on node3</i>	Later in the section	
Non-SAN data LIFs owned by node2	Stage 3: <i>Moving non-SAN data LIFs owned by node2 to node3</i>	Later in the section	
Cluster, intercluster, node-management, cluster-management, and physical ports	Stage 4: <i>Recording node2 information</i>	Stage 5: <i>Installing and booting node4</i> Stage 5: <i>Mapping ports from node2 to node4</i>	
Cluster network ports on node4	Stage 5: <i>Mapping ports from node2 to node4</i>	Later in the section	
Available ports and broadcast domains on node4	Stage 5: <i>Mapping ports from node2 to node4</i>	Later in the section	
Private and public SSL certificates for the storage system and private SSL certificates for each key management server	Stage 6: <i>Setting up Storage Encryption on the new nodes</i>	Later in the section	

## Upgrading the node pair

To upgrade the node pair, you need to prepare the original nodes and then perform a series of steps on both the original and new nodes. You can then decommission the original nodes.

### Steps

1. [Upgrading the node pair: Stage 1](#) on page 11
2. [Upgrading the node pair: Stage 2](#) on page 26
3. [Upgrading the node pair: Stage 3](#) on page 37
4. [Upgrading the node pair: Stage 4](#) on page 61
5. [Upgrading the node pair: Stage 5](#) on page 65
6. [Upgrading the node pair: Stage 6](#) on page 87

## Upgrading the node pair: Stage 1

During Stage 1, you must prepare the nodes for the upgrade and might need to rekey disks for Storage Encryption. You must also prepare to netboot the new controllers if they do not have the same or later version of ONTAP 9.x that is installed on the original controllers.

### Steps

1. [Determining whether the controller has aggregates on internal disk drives](#) on page 11
2. [Preparing the nodes for the upgrade](#) on page 14
3. [Getting an IP address of an external key management server for storage encryption](#) on page 24
4. [Quiesce SnapMirror relationship](#) on page 25
5. [Preparing for netboot](#) on page 25

## Determining whether the controller has aggregates on internal disk drives

If you are upgrading controllers with internal disk drives, you need to complete several commands and examine their output to ensure that none of the internal disk drives contains root aggregates or data aggregates.

### About this task

If you are not upgrading controllers with aggregates on internal disk drives, skip this section and go to the section *Preparing the nodes for the upgrade*.

### Steps

1. Enter the nodeshell, once for each of the original nodes:

```
system node run -node node-name
```

2. Display the internal drives:

```
sysconfig -av
```

### Example

The system displays detailed information about the node's configuration, including storage, as seen in the partial output shown in the following example:

```
node> sysconfig -av
slot 0: SAS Host Adapter 0a (PMC-Sierra PM8001 rev. C, SAS, UP)
  Firmware rev:      01.11.06.00
  Base WWN:          5:00a098:0008a3b:b0
  Phy State:         [0] Enabled, 6.0 Gb/s
                    [1] Enabled, 6.0 Gb/s
                    [2] Enabled, 6.0 Gb/s
                    [3] Enabled, 6.0 Gb/s

  ID      Vendor  Model          FW      Size
  00.0 : NETAPP X306_HMARK02TSSM NA04 1695.4GB (3907029168 512B/sect)
  00.1 : NETAPP X306_HMARK02TSSM NA04 1695.4GB (3907029168 512B/sect)
  00.2 : NETAPP X306_HMARK02TSSM NA04 1695.4GB (3907029168 512B/sect)
  00.3 : NETAPP X306_HMARK02TSSM NA04 1695.4GB (3907029168 512B/sect)
  00.4 : NETAPP X306_HMARK02TSSM NA04 1695.4GB (3907029168 512B/sect)
  00.5 : NETAPP X306_HMARK02TSSM NA04 1695.4GB (3907029168 512B/sect)
  00.6 : NETAPP X306_HMARK02TSSM NA04 1695.4GB (3907029168 512B/sect)
  00.7 : NETAPP X306_HMARK02TSSM NA04 1695.4GB (3907029168 512B/sect)
  00.8 : NETAPP X306_HMARK02TSSM NA04 1695.4GB (3907029168 512B/sect)
  00.9 : NETAPP X306_HMARK02TSSM NA04 1695.4GB (3907029168 512B/sect)
  00.10: NETAPP X306_HMARK02TSSM NA04 1695.4GB (3907029168 512B/sect)
  00.11: NETAPP X306_HMARK02TSSM NA04 1695.4GB (3907029168 512B/sect)
  ...
```

- Examine the storage output of the `sysconfig -av` command to identify the internal disk drives, and then record the information.

Internal drives have 00. at the beginning of their ID. The 00. indicates an internal disk shelf, and the number after the decimal point indicates the individual disk drive.

- Enter the following command on both controllers:

```
aggr status -r
```

#### Example

The system displays aggregate status of the node, as shown in the partial output of the following example:

```
node> aggr status -r
Aggregate aggr2 (online, raid_dp, parity uninit'd!) (block checksums)
Plex /aggr2/plex0 (online, normal, active)
RAID group /aggr2/plex0/rg0 (normal, block checksums)

RAID Disk Device      HA  SHELF  BAY  CHAN  Pool  Type  RPM  Used (MB/blks)  Phys (MB/blks)
-----
dparity 0a.00.1 0a   0    1   SA:B  0   BSAS  7200 1695466/3472315904 1695759/3472914816
parity  0a.00.3 0a   0    3   SA:B  0   BSAS  7200 1695466/3472315904 1695759/3472914816
data    0a.00.9 0a   0    9   SA:B  0   BSAS  7200 1695466/3472315904 1695759/3472914816
...
```

**Note:** The device used to create the aggregate may not be a physical disk but might be a partition.

- Examine the output of the `aggr status -r` command to identify the aggregates using internal disk drives, and then record the information.

In the example in Step 4, `aggr2` uses internal drives, as indicated by the shelf ID of 0.

- Enter the following command on both controllers:

```
aggr status -v
```

#### Example

The system displays information about the volumes on the aggregate, as shown in the partial output in the following example:

```

node> aggr status -v
...
  aggr2      online      raid_dp, aggr      nosnap=off, raidtype=raid_dp,
raidsize=14,
              64-bit      raid_lost_write=on, ignore_inconsistent=off,
rlw_on        snapmirrored=off, resyncsnaptime=60,
              fs_size_fixed=off, lost_write_protect=on,
              ha_policy=cfo, hybrid_enabled=off,
              percent_snapshot_space=0%,
              free_space_realloc=off, raid_cv=on,
              thorough_scrub=off

              Volumes: vol6, vol5, vol14
...
  aggr0      online      raid_dp, aggr      root, diskroot, nosnap=off, raidtype=raid_dp,
raidsize=14, raid_lost_write=on,
              ignore_inconsistent=off, snapmirrored=off,
              resyncsnaptime=60, fs_size_fixed=off,
              lost_write_protect=on, ha_policy=cfo,
              hybrid_enabled=off, percent_snapshot_space=0%,
              free_space_realloc=off, raid_cv=on

              Volumes: vol0

```

Based on the output in Step 4 and Step 6, aggr2 uses three internal drives—0a.00.1, 0a.00.3, and 0a.00.9—and the volumes on aggr2 are vol6, vol5, and vol14. Also, in the output of Step 6, the readout for aggr0 contains the word `root` at the beginning of the information for the aggregate. That indicates that it contains a root volume.

7. Examine the output of the `aggr status -v` command to identify the volumes belonging to any aggregates that are on an internal drive and whether any of those volumes contain a root volume.
8. Exit the nodeshell by entering the following command on each controller:

**exit**

9. Take one of the following actions:

If the controllers...	Then...
Do not contain any aggregates on internal disk drives	Continue with this procedure.
Contain aggregates but no volumes on the internal disk drives	Continue with this procedure. <b>Note:</b> Before you continue, you must place the aggregates offline, and then destroy the aggregates on the internal disk drives. See the <i>ONTAP 9 Logical Storage Management Guide</i> for information about managing aggregates.
Contain non-root volumes on the internal drives	Continue with this procedure. <b>Note:</b> Before you continue, you must move the volumes to an external disk shelf, place the aggregates offline, and then destroy the aggregates on the internal disk drives. See the <i>ONTAP 9 Logical Storage Management Guide</i> for information about moving volumes.
Contain root volumes on the internal drives	Do not continue with this procedure. You can upgrade the controllers by using the procedure <i>Upgrading the controller hardware on a pair of nodes running clustered Data ONTAP by moving volumes</i> on the NetApp Support Site at <a href="http://mysupport.netapp.com">mysupport.netapp.com</a> .
Contain non-root volumes on the internal drives and you cannot move the volumes to external storage	Do not continue with this procedure. You can upgrade the controllers by using the procedure <i>Upgrading the controller hardware on a pair of nodes running clustered Data ONTAP by moving volumes</i> on the NetApp Support Site at <a href="http://mysupport.netapp.com">mysupport.netapp.com</a> .

## Preparing the nodes for the upgrade

Before you can replace the original nodes, you must ensure that they are in an HA pair, have no missing or failed disks, can access each other's storage, and do not own data LIFs assigned to the other nodes in the cluster. You also need to collect information about the original nodes and, if the cluster is in a SAN environment, ensure that all the nodes in the cluster are in quorum.

### Steps

1. Ensure that each of the original nodes has enough resources to adequately support the workload of both nodes during takeover mode.

Follow the *Best practices for HA pairs* section in the *ONTAP 9 High-Availability Configuration Guide*. Neither of the original nodes should be running at more than 50 percent utilization; if a node is running at less than 50 percent utilization, it can handle the loads for both nodes during the controller upgrade.

2. Complete the following substeps to create a performance baseline for the original nodes:

- a. Make sure that the diagnostic user account is unlocked.

**Important:** The diagnostic user account is intended only for low-level diagnostic purposes and should be used only with guidance from technical support.

**Important:** See *System Administration Reference* for information about unlocking the user accounts.

- b. Download the Performance and Statistics Collector (Perfstat Converged) tool from the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com).

Perfstat Converged tool lets you establish a performance baseline for comparison after the upgrade.

- c. Create a performance baseline, following the instructions on the NetApp Support Site.

3. Open a support case on the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com).

You can use the case to report any issues that might arise during the upgrade.

4. Verify that NVMEM or NVRAM batteries of node3 and node4 are charged, and charge them if they are not.

You need to physically check node3 and node4 to see if the NVMEM or NVRAM batteries are charged. Check the *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com) for information about the LEDs for the model of node3 and node4.

You can check the battery status either by connecting to a console or using SSH, logging into the Service Processor (SP) for your system, and use the system sensors to see if the battery has a sufficient charge.

**Attention:** Do not try to clear the NVRAM contents. If there is a need to clear the contents of NVRAM, contact NetApp technical support.

5. Check the version of ONTAP on node3 and node4.

The new nodes must have the same version of ONTAP 9.x installed on them that is installed on the original nodes. If the new nodes have a different version of ONTAP installed, you need to netboot the new controllers after you install them. See the *ONTAP 9 Upgrade and Revert/Downgrade Guide* for instructions to upgrade ONTAP.

Information about the version of ONTAP on node3 and node4 should be included in the shipping boxes. The ONTAP version is displayed when the node boots up or you can boot the node to maintenance mode and run the command:

```
version
```

6. Check whether you have two or four cluster LIFs on node1 and node 2:

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

## Example

The system displays any cluster LIFs, as shown in the following example:

```
cluster::> network interface show -role cluster
Vserver      Logical      Status      Network      Current      Current      Is
-----      -
node1        clus1        up/up       172.17.177.2/24  node1        e0c          true
              clus2        up/up       172.17.177.6/24  node1        e0e          true
node2        clus1        up/up       172.17.177.3/24  node2        e0c          true
              clus2        up/up       172.17.177.7/24  node2        e0e          true
```

7. If you have two or four cluster LIFs on node1 or node2, make sure that you can ping both cluster LIFs across all the available paths by completing the following substeps:

- a. Enter the advanced privilege level:

```
set -privilege advanced
```

The system displays the following message:

```
Warning: These advanced commands are potentially dangerous; use them only when
directed to do so by NetApp personnel.
do you wish to continue? (y or n):
```

- b. Enter **y**.
- c. Ping the nodes and test the connectivity:

```
cluster ping-cluster -node node_name
```

## Example

The system displays a message similar to the following example:

```
cluster::*> cluster ping-cluster -node node1
Host is node1
Getting addresses from network interface table...
Local = 10.254.231.102 10.254.91.42
Remote = 10.254.42.25 10.254.16.228
Ping status:
....
Basic connectivity succeeds on 4 path(s) Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 4 path(s):
Local 10.254.231.102 to Remote 10.254.16.228
Local 10.254.231.102 to Remote 10.254.42.25
Local 10.254.91.42 to Remote 10.254.16.228
Local 10.254.91.42 to Remote 10.254.42.25
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

If the node uses two cluster ports, you should see that it is able to communicate on four paths, as shown in the example.

- d. Return to the administrative level privilege:

```
set -privilege admin
```

8. Ensure that node1 and node2 are in an HA pair and verify that the nodes are connected to each other, and that takeover is possible:

## storage failover show

### Example

The following example shows the output when the nodes are connected to each other and takeover is possible:

```
cluster::> storage failover show
Node           Partner           Takeover
Possible State Description
-----
node1          node2              true      Connected to node2
node2          node1              true      Connected to node1
```

Neither node should be in partial giveback. The following example shows that node1 is in partial giveback:

```
cluster::> storage failover show
Node           Partner           Takeover
Possible State Description
-----
node1          node2              true      Connected to node2, Partial giveback
node2          node1              true      Connected to node1
```

If either node is in partial giveback, use the `storage failover giveback` command to perform the giveback, and then use the `storage failover show-giveback` command to make sure that no aggregates still need to be given back. See the *ONTAP 9 High-Availability Configuration Guide* for detailed information about the commands.

9. Ensure that neither node1 nor node2 owns the aggregates for which it is the current owner (but not the home owner):

```
storage aggregate show -node node_name -is-home false -fields owner-name,home-name,state
```

### Example

If neither node1 nor node2 owns aggregates for which it is the current owner (but not the home owner), the system will return a message similar to the following example:

```
cluster::> storage aggregate show -node node2 -is-home false -fields owner-name,home-
name,state
There are no entries matching your query.
```

The following example shows the output of the command for a node named node2 that is the home owner, but not the current owner, of four aggregates:

```
cluster::> storage aggregate show -node node2 -is-home false -fields owner-name,home-
name,state
aggregate    home-name    owner-name    state
-----
aggr1        node1        node2         online
aggr2        node1        node2         online
aggr3        node1        node2         online
aggr4        node1        node2         online
4 entries were displayed.
```

10. Take one of the following actions:

If the command in Step 9...	Then...
Had blank output	Skip Step 11 and go to Step 12.
Had output	Go to Step 11.

11. If either node1 or node2 owns aggregates for which it is the current owner but not the home owner, complete the following substeps:



- a. Return the aggregates currently owned by the partner node to the home owner node:

```
storage failover giveback -ofnode home_node_name
```

- b. Verify that neither node1 nor node2 still owns aggregates for which it is the current owner (but not the home owner):

```
storage aggregate show -nodes node_name -is-home false -fields owner-name,home-name,state
```

### Example

The following example shows the output of the command when a node is both the current owner and home owner of aggregates:

```
cluster::> storage aggregate show -nodes node1 -is-home true -fields owner-name,home-name,state
aggregate      home-name      owner-name      state
-----
aggr1          node1          node1           online
aggr2          node1          node1           online
aggr3          node1          node1           online
aggr4          node1          node1           online
4 entries were displayed.
```

12. Ensure that node1 and node2 can access each other's storage and verify that no disks are missing:

```
storage failover show -fields local-missing-disks,partner-missing-disks
```

### Example

The following example shows the output when no disks are missing:

```
cluster::> storage failover show -fields local-missing-disks,partner-missing-disks
node      local-missing-disks  partner-missing-disks
-----
node1     None                None
node2     None                None
```

If any disks are missing, see the *ONTAP 9 Disks and Aggregates Power Guide*, the *ONTAP 9 Logical Storage Management Guide*, and the *ONTAP 9 High-Availability Configuration Guide* on the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com) to configure storage for the HA pair.

13. Ensure that node1 and node2 are healthy and eligible to participate in the cluster:

```
cluster show
```

### Example

The following example shows the output when both nodes are eligible and healthy:

```
cluster::> cluster show
Node      Health  Eligibility
-----
node1     true    true
node2     true    true
```

14. Ensure that node1 and node2 are running the same ONTAP release:

```
system node image show -node node1,node2 -iscurrent true
```

### Example

The following example shows the output of the command:

```
cluster::> system node image show -node node1,node2 -iscurrent true
Node      Image      Is      Is      Install
Default  Current  Version  Date
-----
node1     image1    true     true    9.1      2/7/2017 20:22:06
node2     image1    true     true    9.1      2/7/2017 20:20:48
2 entries were displayed.
```

15. Verify that neither node1 nor node2 owns any data LIFs that belong to other nodes in the cluster and check the Current Node and Is Home columns in the output:

```
network interface show -role data -is-home false -curr-node node_name
```

### Example

The following example shows the output when node1 has no LIFs that are home-owned by other nodes in the cluster:

```
cluster::> network interface show -role data -is-home false -curr-node node1
There are no entries matching your query.
```

The following example shows the output when node1 owns data LIFs home-owned by the other node:

```
cluster::> network interface show -role data -is-home false -curr-node node1
Logical      Status      Network      Current      Current Is
Interface    Admin/Oper  Address/Mask Node          Port        Home
-----
vs0          data1       up/up        172.18.103.137/24 node1        e0d         false
             data2       up/up        172.18.103.143/24 node1        e0f         false
2 entries were displayed.
```

16. If the output in Step 15 shows that either node1 or node2 owns any data LIFs home-owned by other nodes in the cluster, migrate the data LIFs away from node1 or node2:

```
network interface revert -vserver * -lif *
```

See the *ONTAP 9 Commands: Manual Page Reference* for detailed information about the `network interface revert` command.

17. Check whether node1 or node2 owns any failed disks:

```
storage disk show -nodelist node1,node2 -broken
```

If any of the disks have failed, remove them, following instructions in the *ONTAP 9 Disks and Aggregates Power Guide*.

18. Collect information about node1 and node2 by completing the following substeps and recording the output of each command:

**Note:** You will use this information later in the procedure.

- a. Record the model, system ID, and serial number of both nodes:

```
system node show -node node1,node2 -instance
```

**Note:** You will use the information to reassign disks and decommission the original nodes.

- b. Enter the following command on both node1 and node2 and record information about the shelves, number of disks in each shelf, flash storage details, memory, NVRAM, and network cards from the output:

```
run -node node_name sysconfig
```

**Note:** You can use the information to identify parts or accessories that you might want to transfer to node3 or node4. If you do not know if the nodes are V-Series systems or have FlexArray Virtualization software, you can learn that also from the output.

- c. Enter the following command on both node1 and node2 and record the aggregates that are online on both nodes:

```
storage aggregate show -node node_name -state online
```

**Note:** You can use this information and the information in the following substep to verify that the aggregates and volumes remain online throughout the procedure, except for the brief period when they are offline during relocation.

- d. Enter the following command on both node1 and node2 and record the volumes that are offline on both nodes:

```
volume show -node node_name -state offline
```

**Note:** After the upgrade, you will run the command again and compare the output with the output in this step to see if any other volumes have gone offline.

19. Enter the following commands to see if any interface groups or VLANs are configured on node1 or node2:

```
network port vlan show
```

```
network port ifgrp show
```

Make note of whether interface groups or VLANs are configured on node1 or node2; you need that information in the next step and later in the procedure.

20. Complete the following substeps on both node1 and node2 to ensure that physical ports can be mapped correctly later in the procedure:

- a. Enter the following command to see if there are failover groups on the node other than `clusterwide`:

```
network interface failover-groups show
```

Failover groups are sets of network ports present on the system. Because upgrading the controller hardware can change the location of physical ports, failover groups can be inadvertently changed during the upgrade.

### Example

The system displays failover groups on the node, as shown in the following example:

```
cluster::> network interface failover-groups show
Vserver          Group          Failover
-----          -
Cluster
                  Cluster
                  node1:e0a, node1:e0b
                  node2:e0a, node2:e0b

fg_6210_e0c
                  Default
                  node1:e0c, node1:e0d
                  node1:e0e, node2:e0c
                  node2:e0d, node2:e0e
2 entries were displayed.
```

- b. If there are failover groups present other than `clusterwide`, record the failover group names and the ports that belong to the failover groups.

- c. Enter the following command to see if there are any VLANs configured on the node:

```
network port vlan show -node node_name
```

VLANs are configured over physical ports. If the physical ports change, then the VLANs will need to be re-created later in the procedure.

### Example

The system displays VLANs configured on the node, as shown in the following example:

```
cluster::> network port vlan show
Network Network
Node   VLAN Name Port   VLAN ID MAC Address
-----
node1  e1b-70  e1b   70     00:15:17:76:7b:69
```

d. If there are VLANs configured on the node, take note of each network port and VLAN ID pairing.

21. Take one of the following actions:

If interface groups or VLANS are...	Description
On node1 or node2	Complete Step 22 and Step 23.
Not on node1 or node2	Go to Step 23.

22. If you do not know if node1 and node2 are in a SAN or non-SAN environment, enter the following command and examine its output:

```
network interface show -vserver vsver_name -data-protocol iscsi|fc
```

### Example

If neither iSCSI nor FC is configured for the SVM, the command will display a message similar to the following example:

```
cluster::> network interface show -vserver vsver_name -data-protocol iscsi|fc
There are no entries matching your query.
```

You can confirm that the node is in a NAS environment by using the command `network interface show` with the `-data-protocol nfs|cifs` parameter.

If either iSCSI or FC is configured for the SVM, the command will display a message similar to the following example:

```
cluster::> network interface show -vserver vsver_name -data-protocol fc
Vserver      Logical Interface   Status Admin/Oper Network Address/Mask Current Node Current Port Current Is Home
-----
vs1          vs1_lif1    up/down    172.17.176.20/24 node1      0d      true
```

23. Verify that all the nodes in the cluster are in quorum by completing the following substeps:

a. Enter the advanced privilege level:

```
set -privilege advanced
```

The system displays the following message:

```
Warning: These advanced commands are potentially dangerous; use them only when
directed to do so by NetApp personnel.
do you wish to continue? (y or n):
```

b. Enter **y**.

c. Verify the cluster service state in the kernel, once for each node:

```
cluster kernel-service show
```

### Example

The system displays a message similar to the following example:

```
cluster::*> cluster kernel-service show
Master          Cluster          Quorum          Availability    Operational
Node            Node             Status           Status          Status
-----
node1           node1            in-quorum       true            operational
                node2            in-quorum       true            operational

2 entries were displayed.
```

Nodes in a cluster are in quorum when a simple majority of nodes are healthy and can communicate with each other. For more information, see the *System Administration Reference*

d. Return to the administrative privilege level:

```
set -privilege admin
```

24. Take one of the following actions:

If the cluster...	Then...
Has SAN configured	Go to Step 25.
Does not have SAN configured	Go to Step 28.

25. Verify that there are SAN LIFs on node1 and node2 for each SVM that has either SAN iSCSI or FC service enabled by entering the following command and examining its output:

```
network interface show -data-protocol iscsi|fcp -home-node node_name
```

### Example

The command displays SAN LIF information for node1 and node2. The following examples show the status in the Status Admin/Oper column as up/up, indicating that SAN iSCSI and FC service are enabled:

```
cluster::*> network interface show -data-protocol fcp|iscsi
Logical      Status      Network      Current      Current      Is
Vserver      Interface  Admin/Oper  Address/Mask Node          Port         Home
-----
a_vs_iscsi
  data1      up/up      10.228.32.190/21  node1        e0a          true
  data2      up/up      10.228.32.192/21  node2        e0a          true
b_vs_fcp
  data1      up/up      20:09:00:a0:98:19:9f:b0  node1        0c           true
  data2      up/up      20:0a:00:a0:98:19:9f:b0  node2        0c           true
c_vs_iscsi_fcp
  data1      up/up      20:0d:00:a0:98:19:9f:b0  node2        0c           true
  data2      up/up      20:0e:00:a0:98:19:9f:b0  node2        0c           true
  data3      up/up      10.228.34.190/21  node2        e0b          true
  data4      up/up      10.228.34.192/21  node2        e0b          true
```

Alternatively, you can view more detailed LIF information by entering the command `network interface show - instance -data-protocol iscsi|fc`.

26. Capture the default configuration of any FC ports on the original nodes by entering the following command and recording the output for your systems:

**ucadmin show**

**Example**

The command displays information about all FC ports in the cluster, as shown in the following example:

```
cluster::> ucadmin show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
node1	0a	fc	initiator	-	-	online
node1	0b	fc	initiator	-	-	online
node1	0c	fc	initiator	-	-	online
node1	0d	fc	initiator	-	-	online
node2	0a	fc	initiator	-	-	online
node2	0b	fc	initiator	-	-	online
node2	0c	fc	initiator	-	-	online
node2	0d	fc	initiator	-	-	online

8 entries were displayed.

You can use the information after the upgrade to set the configuration of FC ports on the new nodes.

27. If you are upgrading a V-Series system or a system with FlexArray Virtualization software, capture information about the topology of the original nodes by entering the following command and recording the output:

**storage array config show -switch**

**Example**

The system displays topology information, as show in the following example:

```
cluster::> storage array config show -switch
```

Node	LUN Group	LUN Count	Array Name	Array Target Port	Target Side Switch Port	Initiator Side Switch Port	Initiator
node1	0	50	IBM_1818FAStT_1	205700a0b84772da	vgbr6510s164:5	vgbr6510s164:3	
				0d			
				206700a0b84772da	vgbr6510s164:6	vgbr6510s164:4	2b
				207600a0b84772da	vgbr6510s163:6	vgbr6510s163:1	0c
node2	0	50	IBM_1818FAStT_1	205700a0b84772da	vgbr6510s164:5	vgbr6510s164:1	0d
				206700a0b84772da	vgbr6510s164:6	vgbr6510s164:2	2b
				207600a0b84772da	vgbr6510s163:6	vgbr6510s163:3	0c
				208600a0b84772da	vgbr6510s163:5	vgbr6510s163:4	2a

7 entries were displayed.

28. Complete the following substeps:
  - a. Enter the following command on one of the original nodes and record the output:

**service-processor show -node \* -instance**

The system displays detailed information about the SP on both nodes.

- b. Ensure that the SP status is online.
  - c. Ensure that the SP network is configured.
  - d. Record the IP address and other information about the SP.

You might want to reuse the network parameters of the remote management devices, in this case the SPs, from the original system for the SPs on the new nodes.

For detailed information about the SP, see the *System Administration Reference* and the *ONTAP 9 Commands: Manual Page Reference*.

29. If you want the new nodes to have the same licensed functionality as the original nodes, enter the following command to see the cluster licenses on the original system:

```
system license show -owner *
```

### Example

The following example shows the site licenses for cluster1:

```
system license show -owner *
Serial Number: 1-80-000013
Owner: cluster1
Package           Type      Description           Expiration
-----
Base              site     Cluster Base License -
NFS               site     NFS License          -
CIFS              site     CIFS License         -
SnapMirror        site     SnapMirror License   -
FlexClone         site     FlexClone License    -
SnapVault         site     SnapVault License    -
6 entries were displayed.
```

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

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

31. Check whether the original system has AutoSupport enabled by entering the following command on each node and examining its output:

```
system node autosupport show -node node1,node2
```

### Example

The command output shows whether AutoSupport is enabled, as shown in the following example:

```
cluster::> system node autosupport show -node node1,node2
Node           State      From           To           Mail Hosts
-----
node1          enable    Postmaster     admin@netapp.com
              mailhost
node2          enable    Postmaster     -           mailhost
2 entries were displayed.
```

32. Take one of the following actions:

If the original system...	Then...
Has AutoSupport enabled...	<ul style="list-style-type: none"> <li>a. Go to Step 33.</li> <li>b. Go to the section <i>Rekeying disks for Storage Encryption</i>.</li> </ul>
Does not have AutoSupport enabled...	<ul style="list-style-type: none"> <li>a. Enable AutoSupport, following the instructions in the <i>System Administration Reference</i>. <ul style="list-style-type: none"> <li><b>Note:</b> AutoSupport is enabled by default when you configure your storage system for the first time. Although you can disable AutoSupport at any time, you should leave it enabled. Enabling AutoSupport can significantly help identify problems and solutions should a problem occur on your storage system.</li> </ul> </li> <li>b. Go to the <i>Rekeying disks for Storage Encryption</i> section.</li> </ul>

33. Verify that AutoSupport is configured with the correct mailhost details and recipient e-mail IDs by entering the following command on both of the original nodes and examining the output:

```
system node autosupport show -node node_name -instance
```

For detailed information about AutoSupport, see the *System Administration Reference* and the *ONTAP 9 Commands: Manual Page Reference*.

34. Send an AutoSupport message to NetApp for node1 by entering the following command:

```
system node autosupport invoke -node node1 -type all -message "Upgrading node1 from platform_old to platform_new"
```

#### Example

```
cluster::> system node autosupport invoke -node node1 -type all -message "Upgrading node1 from FAS6240 to FAS6290"
```

**Note:** Do not send an AutoSupport message to NetApp for node2 at this point; you do so later in the procedure.

35. Verify that the AutoSupport message was sent by entering the following command and examining its output:

```
system node autosupport show -node node1 -instance
```

The fields `Last Subject Sent:` and `Last Time Sent:` contain the message title of the last message sent and the time the message was sent.

## Getting an IP address of an external key management server for storage encryption

Beginning with clustered Data ONTAP 8.3.1, you manage the Storage Encryption functionality in Data ONTAP using clustershell CLI commands instead of individual nodeshell commands. This includes managing authentication keys. After upgrading, you must immediately configure Storage Encryption and establish a cluster-wide authentication key to replace the previous node-level authentication keys.

### Steps

1. Install the necessary client and server SSL certificates required to communicate with key management servers by using the `security certificate install` command.
2. Configure Storage Encryption on all nodes by using the `security key-manager setup` command on each node.
3. Add the IP address for each key management server by using the `security key-manager add` command.
4. Verify that the same key management servers are configured and available on all nodes in the cluster:

```
security key-manager show -status
```

5. Create a new cluster-wide authentication key:

```
security key-manager create-key
```

6. Make a note of the new authentication key ID.

7. Rekey all SEDs with the new authentication key:

```
storage encryption disk modify -disk * -data-key-id authentication_key_id
```



## Quiesce SnapMirror relationship

Before you netboot the system, you must ensure that all the SnapMirror relationships are quiesced. When a SnapMirror relationship is quiesced, it remains quiesced across reboots and fail-overs.

### Steps

1. Verify the SnapMirror relationship state on the destination cluster:

```
snapmirror show
```

**Note:** If the state is `Transferring`, you must abort those transfers by using the

```
snapmirror abort -destination vserver
```

command. The abort fails if the SnapMirror relationship is not in the `Transferring` state.

2. Quiesce all relationships between the cluster:

```
snapmirror quiesce-destination vserver
```

## Preparing for netboot

After you install new nodes, you might need to netboot them. The term *netboot* means you are booting from a Data ONTAP image stored on a remote server. When preparing for netboot, you must put a copy of the ONTAP 9.x boot image on a Web server that the system can access.

### Before you begin

- You must have verified that you can access an HTTP server with the system before and after installing node3 and node4.
- You must have verified that you have access to the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com). You need to download the necessary system files for your platform and the correct version of ONTAP from the NetApp Support Site.

### About this task

You must netboot the new controllers if they do not have the same or later version of ONTAP 9 installed on them that is installed on the original controllers. After you install each new controller, you boot the system from the ONTAP 9 image stored on the Web server. You can then download the correct files to the boot media device for subsequent system boots.

However, you do not need to netboot the controllers if the same or later version of ONTAP 9 is installed on them that is installed on the original controllers. If so, you can skip this section and proceed to *Upgrading the node pair: Stage 2*.

### Steps

1. Download and extract the files used for performing the netboot of the system:
  - a. Download the appropriate `netboot.tgz` file from the software download section of the NetApp Support Site to a web-accessible directory.
  - b. Change to the web-accessible directory.
  - c. Extract the contents of the `netboot.tgz` file to the target directory:

```
tar -zxvf netboot.tgz
```

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

Your directory listing should contain the following directory:

```
netboot/
```

2. Download the `image.tgz` file from the NetApp Support Site to the web-accessible directory that you chose in Step 1a. Your directory listing should contain the following files and directory:

```
image.tgz
netboot.tgz
netboot/
```

You will use information in the directories in Stage 3.

## Upgrading the node pair: Stage 2

During Stage 2, you relocate the node1 aggregates and LIFs to node2, record node1 information, and then retire node1.

### Steps

1. [Relocating non-root aggregates from node1 to node2](#) on page 26
2. [Moving non-SAN data LIFs owned by node1 to node2](#) on page 29
3. [Recording node1 information](#) on page 32
4. [Retiring node1](#) on page 34

## Relocating non-root aggregates from node1 to node2

Before you can replace node1 with node3, you need to move the non-root aggregates from node1 to node2 by using the `storage aggregate relocation` command and then verifying the relocation.

### Steps

1. Relocate the non-root aggregates by completing the following substeps:
  - a. Access the advanced privilege level by entering the following command on either node:

```
set -privilege advanced
```

The system displays the following message:

```
Warning: These advanced commands are potentially dangerous; use them only when
directed to do so by NetApp personnel.
do you wish to continue? (y or n):
```

- b. Enter `y`.

- c. Enter the following command:

```
storage aggregate relocation start -node node1 -destination node2 -aggregate-list * -ndo-
controller-upgrade true
```

The system displays the following message:

```
Warning: Aggregate relocation will not change home ownership of an aggregate which is
owned by the source node during an NDO controller upgrade. This parameter should be
used only while performing NDO controller upgrade.
Are you performing NDO controller upgrade? {y|n}:
```

- d. When prompted, enter `y`.

Relocation will occur in the background. It could take anywhere from a few seconds to a couple of minutes to relocate an aggregate. The time includes both client outage and non-outage portions. The command does not relocate any offline or restricted aggregates.

e. Return to the admin level by entering the following command:

```
set -privilege admin
```

2. Check the relocation status by entering the following command on node1:

```
storage aggregate relocation show -node node1
```

The output will display Done for an aggregate after it has been relocated.

**Note:** Wait until all non-root aggregates owned by node1 have been relocated to node2 before proceeding to the next step.

3. Take one of the following actions:

<b>If relocation...</b>	<b>Then..</b>
Of all aggregates is successful	Go to Step 4.

If relocation...	Then..
Of any aggregates fails or is vetoed	<p>a. Check the EMS logs for the corrective action.</p> <p>b. Perform the corrective action.</p> <p>c. Access the advanced privilege level by entering the following command on either node:  <b>set -privilege advanced</b>            The system displays the following message:</p> <pre style="background-color: #f0f0f0; padding: 5px;">Warning: These advanced commands are potentially dangerous; use them only when directed to do so by NetApp personnel. do you wish to continue? {y or n}:</pre> <p>d. Enter <b>y</b>.</p> <p>e. Relocate any failed or vetoed aggregates by entering the following command:  <b>storage aggregate relocation start -node node1 -destination node2 -aggregate-list * -ndo-controller-upgrade true</b>            The system displays the following message:</p> <pre style="background-color: #f0f0f0; padding: 5px;">Warning: Aggregate relocation will not change home ownership of an aggregate which is owned by the source node during an NDO controller upgrade. This parameter should be used only while performing NDO controller upgrade. Are you performing NDO controller upgrade? {y n}:</pre> <p>f. When prompted, enter <b>y</b>.</p> <p>g. Return to the admin level by entering the following command:  <b>set -privilege admin</b></p> <p>If necessary, you can force the relocation using one of the following methods:</p> <ul style="list-style-type: none"> <li>• Overriding veto checks by using the command <code>storage aggregate relocation start -override-vetoes true -ndo-controller-upgrade</code></li> <li>• Overriding destination checks by using the command <code>storage aggregate relocation start -override-destination-checks true -ndo-controller-upgrade</code></li> </ul> <p>See the <i>ONTAP 9 Disks and Aggregates Power Guide</i> and the <i>ONTAP 9 Commands: Manual Page Reference</i> for more information about storage aggregate relocation commands.</p>

4. Verify that all the non-root aggregates are online and their state on node2 by entering the following command:

```
storage aggregate show -node node2 -state online -root false
```

### Example

The following example shows that the non-root aggregates on node2 are online:

```
cluster::> storage aggregate show -node node2 state online -root false
Aggregate      Size Available Used% State  #Vols  Nodes      RAID Status
-----
aggr_1
           744.9GB   744.8GB   0% online    5 node2      raid_dp,
```

```

aggr_2      825.0GB   825.0GB   0% online      1 node2      normal
            raid_dp,
            normal
2 entries were displayed.

```

If the aggregates have gone offline or become foreign on node2, bring them online by using the following command on node2, once for each aggregate:

```
storage aggregate online -aggregate aggr_name
```

- Verify that all the volumes are online on node2 by entering the following command on node2 and examining its output:

```
volume show -node node2 -state offline
```

If any volumes are offline on node2, bring them online by using the following command on node2, once for each volume:

```
volume online -vserver vservice-name -volume volume-name
```

The *vservice-name* to use with this command is found in the output of the previous `volume show` command.

- Enter the following command on node2:

```
storage failover show -node node2
```

The output should display the following message:

```
Node owns partner's aggregates as part of the nondisruptive controller upgrade procedure.
```

- Verify that node1 does not own any non-root aggregates that are online by entering the following command:

```
storage aggregate show -owner-name node1 -ha-policy sfo -state online
```

The output should not display any online non-root aggregates, which have already been relocated to node2.

## Moving non-SAN data LIFs owned by node1 to node2

Before you can replace node1 with node3, you need to move the non-SAN data LIFs owned by node1 to node2 if you have a two-node cluster, or to a third node if your cluster has more than two nodes. The method you use depends on whether the cluster is configured for NAS or SAN.

### About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. Moving SAN LIFs is not necessary for cluster or service health during the upgrade. You must verify that the LIFs are healthy and located on appropriate ports after you bring node3 online.

### Steps

- List all the non-SAN data LIFs owned by node1 by entering the following command on either node and capturing the output:

```
network interface show -data-protocol nfs|cifs -curr-node node1
```

### Example

The system displays the non-SAN data LIFs on node1 as shown in the following example:

```

cluster::> network interface show -data-protocol nfs|cifs -curr-node node1
Logical      Status      Network      Current      Current      Is
Vserver      Interface   Admin/Oper   Address/Mask Node          Port         Home
-----
vs0
a0a          up/down    10.63.0.53/24 node1        a0a          true
data1       up/up      10.63.0.50/18 node1        e0c          true
rads1       up/up      10.63.0.51/18 node1        e1a          true

```

vs1	rads2	up/down	10.63.0.52/24	node1	e1b	true
	lif1	up/up	192.17.176.120/24	node1	e0c	true
	lif2	up/up	172.17.176.121/24	node1	e1a	true

2. Take one of the following actions:

If node1...	Description
Has interface groups or VLANs configured	Go to Step 3.
Does not have interface groups or VLANs configured	Skip Step 3 and go to Step 4.

Use the `network port vlan show` to display information about the network ports attached to VLANs, and use the `network port ifgrp show` to display information about the port interface groups.

3. Take the following steps to migrate any non-SAN data LIFs hosted on interface groups and VLANs on node1:

a. Migrate the LIFs hosted on any interface groups and the VLANs on node1 to a port on node2 that is capable of hosting LIFs on the same network as that of the interface groups by entering the following command, once for each LIF:

```
network interface migrate -vserver vsserver_name -lif lif_name -destination-node node2 -destination-port netport|ifgrp
```

b. Modify the home port and the home node of the LIFs and VLANs in substep a to the port and node currently hosting the LIFs by entering the following command, once for each LIF:

```
network interface modify -vserver vsserver_name -lif lif_name -home-node node2 -home-port netport|ifgrp
```

4. Take one of the following actions:

If the cluster is configured for...	Then...
NAS	Complete Step 5 through Step 8.
SAN	Disable all the SAN LIFs on the node to take them down for the upgrade:  <pre>network interface modify -vserver vsserver_name -lif lif_name -home-node node_to_upgrade -home-port netport ifgrp -status-admin down</pre>

5. Migrate non-SAN data LIFs from node1 to node2 by entering the following command, once for each data LIF:

```
network interface migrate -vserver vsserver_name -lif lif_name -destination-node node2 -destination-port data_port
```

6. Enter the following command and examine its output to verify that LIFs have been moved to the correct ports and that the LIFs have the status of up by entering the following command on either node and examining the output:

```
network interface show -curr-node node2 -data-protocol nfs|cifs
```

7. Enter the following command to modify the home-node of the migrated LIFs:

```
network interface modify -vserver vsserver_name -lif LIF_name -home-node node2 -home-port port_name
```

8. Verify whether the LIF is using the port as its home or current port. If the port is not home or current port then go to step 9:

```
network interface show -home-node node2 -home-port port_name
```

```
network interface show -curr-node node_name -curr-port port_name
```

9. If the LIFs are using the port as a home port or current port, then modify the LIF to use a different port:

```
network interface migrate -vserver vservice -LIF LIF_name -destination-node node_name -
destination-port port_name
```

```
network interface modify -vserver vservice -LIF LIF_name -home-node node_name -home-port
port_name
```

10. If the ports currently hosting data LIFs are not going to exist on the new hardware, remove them from the broadcast domain now:

```
network port broadcast-domain remove-ports -ip-space Default -broadcast-domain Default -ports
node:port
```

11. 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 nodename -status-
admin up
```

**Note:** For MetroCluster configuration, you might not be able to change the broadcast domain of a port because it is associated with a port hosting the LIF of a destination SVM. Enter the following command from the corresponding source SVM on the remote site to reallocate the destination LIF to an appropriate port:

```
metrocluster vservice resync -vserver vservice_name
```

12. Enter the following command and examine its output to verify that there are no data LIFs remaining on node1:

```
network interface show -curr-node node1 -role data
```

13. If you have interface groups or VLANs configured, complete the following substeps:

- a. Remove the VLANs from the interface groups by entering the following command:

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

- b. Enter the following command and examine its output to see if there are any interface groups configured on the node:

```
network port ifgrp show -node nodename -ifgrp ifgrp_name -instance
```

### Example

The system displays interface group information for the node as shown in the following example:

```
cluster::> network port ifgrp show -node node1 -ifgrp a0a -instance
      Node: node1
      Interface Group Name: a0a
      Distribution Function: ip
      Create Policy: multimode_lacp
      MAC Address: 02:a0:98:17:dc:d4
      Port Participation: partial
      Network Ports: e2c, e2d
      Up Ports: e2c
      Down Ports: e2d
```

- c. If any interface groups are configured on the node, record the names of those groups and the ports assigned to them, and then delete the ports by entering the following command, once for each port:

```
network port ifgrp remove-port -node nodename -ifgrp ifgrp_name -port netport
```

## Recording node1 information

Before you can shut down and retire node1, you need to record information about its cluster network, management, and FC ports as well as its NVRAM System ID. You need that information later in the procedure when you map node1 to node3 and reassign disks.

### Steps

1. Enter the following command and capture its output:

```
network route show
```

### Example

The system displays output similar to the following example:

```
cluster::> network route show
```

Vserver	Destination	Gateway	Metric
-----	-----	-----	-----
iscsi_vserver	0.0.0.0/0	10.10.50.1	20
node1	0.0.0.0/0	10.10.20.1	10
.....			
node_1_2	0.0.0.0/0	192.168.1.1	20

2. Enter the following command and capture its output:

```
vserver services name-service dns show
```

### Example

The system displays output similar to the following example:

```
cluster::> vserver services name-service dns show
```

Vserver	State	Domains	Name Servers
-----	-----	-----	-----
node_1_2	enabled	alpha.beta.gamma.netapp.com	10.10.60.10, 10.10.60.20
vs_base1	enabled	alpha.beta.gamma.netapp.com, beta.gamma.netapp.com, gamma.netapp.com,	10.10.60.10, 10.10.60.20
....			
....			
vs_peer1	enabled	alpha.beta.gamma.netapp.com, gamma.netapp.com,	10.10.60.10, 10.10.60.20

3. Find the cluster network and node-management ports on node1 by entering the following command on either controller:

```
network interface show -curr-node node1 -role cluster,intercluster,node-mgmt,cluster-mgmt
```

### Example

The system displays the cluster, intercluster, node-management, and cluster-management LIFs for the node in the cluster, as shown in the following example:

```
cluster::> network interface show -curr-node node1 -role cluster,intercluster,node-  
mgmt,cluster-mgmt
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
-----	-----	-----	-----	-----	-----	-----
vserver1						



```

node1      cluster_mgmt up/up    192.168.0.xxx/24  node1      e0c      true
           intercluster up/up    192.168.1.201/24 node1      e0e      true
           clus1       up/up    169.254.xx.xx/24  node1      e0a      true
           clus2       up/up    169.254.xx.xx/24  node1      e0b      true
           mgmt1      up/up    192.168.0.xxx/24  node1      e0c      true
5 entries were displayed.

```

**Note:** Your system might not have intercluster LIFs.

4. Capture the information in the output of the command in Step 3 to use in the section *Mapping cluster and node-management ports from node1 to node3*.

The output information is required to map the new controller ports to the old controller ports.

5. Enter the following command on node1:

```
network port show -node node1 -type physical
```

### Example

The system displays the physical ports on the node as shown in the following example:

```

cluster::> network port show -node node1 -type physical
Node  Port      IPspace      Broadcast Domain Link  MTU  Speed (Mbps)
-----
node1  e0M      Default      172.17.178.19/24 up    1500 auto/100
       e0a      Default      172.17.178.19/24 up    1500 auto/1000
       e0b      Default      -            up    1500 auto/1000
       e1a      Cluster      Cluster      up    9000 auto/10000
       e1b      Cluster      Cluster      up    9000 auto/1000
5 entries were displayed.

```

6. Record the ports and their broadcast domains.

The broadcast domains will need to be mapped to the new ports on the new controller later in the procedure.

7. Enter the following command on node1:

```
network fcp adapter show -node node1
```

### Example

The system displays the FC ports on the node, as shown in the following example:

```

cluster::> fcp adapter show -node node1
Node      Adapter  Connection Established Host Port Address
-----
node1     0a       ptp       11400
node1     0c       ptp       11700
node1     6a       loop      0
node1     6b       loop      0
4 entries were displayed.

```

8. Record the ports.

The output information is required to map the new FC ports on the new controller later in the procedure.

9. If you did not do so earlier, check whether there are interface groups or VLANs configured on node1 by entering the following commands:

```
network port ifgrp show
```

```
network port vlan show
```

You will use the information in the section *Mapping ports from node1 to node3*.

10. Take one of the following actions:

If you...	Then...
Recorded the NVRAM System ID number in the section <i>Preparing the nodes for the upgrade</i>	Go on to the next section, <i>Retiring node1</i> .
Did not record the NVRAM System ID number in the section <i>Preparing the nodes for the upgrade</i>	Complete Step 11 and Step 12 and then go on to the next section, <i>Retiring node1</i> .

11. Enter the following command on either controller:

```
system node show -instance -node node1
```

### Example

The system displays information about node1 as shown in the following example:

```
cluster::> system node show -instance -node node1

                Node: node1
                Owner:
                Location: GD1
                Model: FAS6240
                Serial Number: 700000484678
                Asset Tag: -
                Uptime: 20 days 00:07
                NVRAM System ID: 1873757983
                System ID: 1873757983
                Vendor: NetApp
                Health: true
                Eligibility: true
```

12. Record the NVRAM System ID number to use in the section *Installing and booting node3*.

## Retiring node1

To retire node1, you need to disable the HA pair with node2, shut node1 down properly, and remove it from the rack or chassis.

### Steps

1. Verify the number of nodes in the cluster:

```
cluster show
```

### Example

The system displays the nodes in the cluster, as shown in the following example:

```
cluster::> cluster show
Node           Health Eligibility
-----
node1          true   true
node2          true   true
2 entries were displayed.
```

2. Disable storage failover, as applicable:

If the cluster is...	Description
A two-node cluster	<p>a. Disable cluster high availability by entering the following command on either node:  <b>cluster ha modify -configured false</b></p> <p>b. Disable storage failover:  <b>storage failover modify -node node1 -enabled false</b></p>
A cluster with more than two nodes	<p>Disable storage failover:  <b>storage failover modify -node node1 -enabled false</b></p>

3. Verify that storage failover was disabled:

```
storage failover show
```

**Example**

The following example shows the output of the `storage failover show` command when storage failover has been disabled for a node:

```
cluster::> storage failover show
Node           Partner      Takeover
Possible State Description
-----
node1          node2        false      Connected to node2, Takeover
is not possible: Storage failover is
disabled

node2          node1        false      Node owns partner's aggregates as part
of the nondisruptive controller upgrade
procedure. Takeover is not
possible: Storage failover is
disabled

2 entries were displayed.
```

4. Verify the data LIF status:

```
network interface show -role data -curr-node node2 -home-node node1
```

Look in the `Status Admin/Oper` column to see if any LIFs are down. If any LIFs are down, consult the “Troubleshooting” section.

5. Take one of the following actions:

If the cluster is...	Description
A two-node cluster	Go to Step 8.
A cluster with more than two nodes	Go to Step 15.

6. Access the advanced privilege level on either node:

```
set -privilege advanced
```

The system displays the following message:

```
Warning: These advanced commands are potentially dangerous; use them only when directed
to do so by NetApp personnel.
do you wish to continue? (y or n):
```

7. Enter **y**.

8. Verify that the cluster HA has been disabled:

```
cluster ha show
```

The system displays the following message:

```
High Availability Configured: false
```

If cluster HA has not been disabled, repeat Step 2.

9. Check whether node1 currently holds epsilon:

```
cluster show
```

Because there is the possibility of a tie in a cluster that has an even number of nodes, one node has an extra fractional voting weight called *epsilon*. See the *System Administration Reference* for more information.

The following example shows that node1 holds epsilon:

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

10. If node1 holds epsilon, then mark epsilon **false** on the node so that it can be transferred to the node2:

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

11. Transfer epsilon to node2 by marking epsilon **true** on node2:

```
cluster modify -node node2 -epsilon true
```

12. Verify that the change to node2 occurred:

```
cluster show
```

### Example

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

The epsilon for node2 should now be **true** and the epsilon for node1 should be **false**.

13. Verify whether the setup is a two-node switchless cluster:

```
network options switchless-cluster show
```

## Example

```
cluster::*> network options switchless-cluster show
Enable Switchless Cluster: false/true
```

The value of this command must match the physical state of the system.

14. Return to the admin level:

```
set -privilege admin
```

15. Halt node1 from the node1 prompt:

```
system node halt -node node1
```

**Attention:** If node1 is in same chassis as node2, do not power off the chassis by using the power switch or by pulling the power cable. If you do so, node2, which is serving data, will go down.

16. When the system prompts you to confirm that you want to halt the system, enter **y**.

The node stops at the boot environment prompt.

17. When node1 displays the boot environment prompt, remove it from the chassis or the rack.

You can decommission node1 after the upgrade is completed. See *Decommissioning the old system*.

## Upgrading the node pair: Stage 3

During Stage 3, you install and boot node3, map the cluster and node-management ports from node1 to node3, verify the node3 installation, and move data LIFs and SAN LIFs belonging to node1 from node2 to node3. You also relocate all aggregates from node2 to node3, and move the data LIFs and SAN LIFs owned by node2 to node3.

### Steps

1. [Installing and booting node3](#) on page 37
2. [Setting the FC or CNA configuration on node3](#) on page 44
3. [Mapping ports from node1 to node3](#) on page 50
4. [Verifying the node3 installation](#) on page 54
5. [Moving non-SAN data LIFs owned by node1 from node2 to node3 and verifying SAN LIFs on node3](#) on page 54
6. [Relocating non-root aggregates from node2 to node3](#) on page 57
7. [Moving non-SAN data LIFs owned by node2 to node3](#) on page 59

## Installing and booting node3

You need to install node3 in the rack, transfer node1's connections to node3, boot node3, and install ONTAP. You then need to reassign any of node1's spare disks, any disks belonging to the root volume, and any non-root aggregates not relocated to node2 earlier.

### About this task

You need to netboot node3 if it does not have the same or a later version of Data ONTAP 8.3 that is installed on node1. After you install node3, boot it from the Data ONTAP 8.3 image stored on the Web server. You can then download the correct files to the boot media device for subsequent system boots.

However, you do not need to netboot node3 if it has the same or a later version of Data ONTAP 8.3 that is installed on node1.

**Important:** If you are upgrading a V-Series system connected to storage arrays or a system with FlexArray Virtualization software that is connected to storage arrays, you need to complete Step 1 through Step 5, leave this section at Step 6 and follow instructions in the “Configuring Fibre Channel ports on node3” and “Checking and configuring CNA ports on node3” sections as needed, entering commands in Maintenance mode. You must then return to this section and resume with Step 7.

However, if you are upgrading a system with storage disks, you need to complete this entire section and then go to the “Configuring Fibre Channel ports on node3” and “Checking and configuring CNA ports on node3” sections, entering commands at the cluster prompt.

## Steps

1. Make sure that you have rack space for node3.

If node1 and node2 were in separate chassis, you can put node3 in the same rack location as node1. However, if node1 was in the same chassis with node2, then you need to put node3 into its own rack space, preferably close to the location of node1.

2. Install node3 in the rack, following instructions in the *Installation and Setup Instructions* for your node model.

**Note:** If you are upgrading to a system with both nodes in the same chassis, install node4 in the chassis as well as node3. If you do not, when you boot node3, the node will behave as if it were in a dual-chassis configuration, and when you boot node4, the interconnect between the nodes will not come up.

3. Cable node3, moving the connections from node1 to node3.

Cable the following connections, using the instructions in the *Installation and Setup Instructions* or the *FlexArray Virtualization Installation Requirements and Reference* for the node3 platform, the appropriate disk shelf guide, and the *ONTAP 9 High-Availability Configuration Guide*:

- Console (remote management port)
- Cluster ports
- Data ports
- Cluster and node management ports
- Storage
- SAN configurations: iSCSI Ethernet and FC switch ports

**Note:** You do not need to move the interconnect card or the cluster interconnect cable connection from node1 to node3 because most platform models have a unique interconnect card model. For the MetroCluster configuration, you need to move the FCVI cable connections from node1 to node3.

4. Turn on the power to node3, and then interrupt the boot process by pressing Ctrl-C at the console terminal to access the boot environment prompt.

If you are upgrading to a system with both nodes in the same chassis, node4 also reboots. However, you can disregard the node4 boot until Step 13.

**Note:** When you boot node3, you might see the following message:

```
WARNING: The battery is unfit to retain data during a power
         outage. This is likely because the battery is
         discharged but could be due to other temporary
         conditions.
         When the battery is ready, the boot process will
         complete and services will be engaged.
         To override this delay, press 'c' followed by 'Enter'
```

5. If you see the warning message in Step 4, take the following actions:

- a. Check for any console messages that might indicate a problem other than a low NVRAM battery, and, if necessary, take any required corrective action.
- b. Allow the battery to charge and the boot process to complete.

**Attention:** Do not override the delay; failure to allow the battery to charge could result in a loss of data.

6. Take one of the following actions:

If your system...	Description
Has disks and no back-end storage	Skip Step 7 through Step 12 and go to Step 13.
Is a V-Series system or a system with FlexArray Virtualization software connected to storage arrays	<p>a. Go to the section “Setting the FC or CNA configuration on node3” and complete the subsections in it, “Setting the onboard Fibre Channel configuration on node3” and “Configuring CNA ports on node3”, as appropriate to your system.</p> <p>b. Return to this section and complete the remaining steps, beginning with Step 7.</p> <p><b>Important:</b> You must reconfigure FC onboard ports, CNA onboard ports, and CNA cards before you boot Data ONTAP on the V-Series or system with FlexArray Virtualization software.</p>

7. Add the FC initiator ports of the new node to the switch zones.

If your system has a tape SAN, then you need zoning for the initiators. See your storage array and zoning documentation for instructions.

8. Add the FC initiator ports to the storage array as new hosts, mapping the array LUNs to the new hosts.

See your storage array and zoning documentation for instructions.

9. Modify the World Wide Port Name (WWPN) values in the host or volume groups associated with array LUNs on the storage array.

Installing a new controller module changes the WWPN and World Wide Node Name (WWNN) values associated with each onboard FC port.

10. If your configuration uses switch-based zoning, adjust the zoning to reflect the new WWNN values.

11. Verify that the array LUNs are now visible to node3:

```
sysconfig -v
```

The system displays all the array LUNs visible to each of the FC initiator ports. If the array LUNs are not visible, you will not be able to reassign disks from node1 to node3 later in this section.

12. At the Maintenance mode prompt, enter the following command:

```
halt
```

The system stops at the boot environment prompt.

13. Take one of the following actions:

If the system you are upgrading to is in a...	Then...
Dual-chassis configuration (with controllers in different chassis)	Go to Step 14.

---

If the system you are upgrading to is in a...	Then...
---	---------

---

Single-chassis configuration (with controllers in the same chassis)

- a. Switch the console cable from node3 to node4.
  - b. Turn on the power to node4, and then interrupt the boot process by pressing Ctrl-C at the console terminal to access the boot environment prompt.  
The power should already be on if both controllers are in the same chassis.  
**Note:** Leave node4 at the boot environment prompt; you will return to node4 in the section “Installing and booting node4”.
  - c. If you see the warning message displayed in Step 4, follow the instructions in Step 5.
  - d. Switch the console cable back from node4 to node3.
  - e. Go to Step 14.
- 

14. Configure node3 for ONTAP:

```
set-defaults
```

15. If FDE is in use on this configuration, the `setenv bootarg.storageencryption.support` command must be set to **true**, and the `kmip.init.maxwait` variable needs to be set to **off** to avoid a boot loop after the node1 configuration is loaded:

```
setenv bootarg.storageencryption.support true
setenv kmip.init.maxwait off
```

16. If the version of Data ONTAP installed on node3 is the same or later than the version of Data ONTAP 8.3 installed on node1, enter the following command:

```
boot_ontap menu
```

**Warning:** If this new node has ever been used in any other cluster or HA pair, you must run `wipeconfig` before proceeding. Failure to do so may result in service outages or data loss. Contact technical support if the replacement controller was previously used, especially if the controllers were running Data ONTAP in 7-Mode.

17. Take one of the following actions:

---

If the system you are upgrading...	Then...
Does <i>not</i> have the correct or current Data ONTAP version on node3	Go to Step 18.
Has the correct or current version of Data ONTAP on node3	Go to Step 23.

---

18. Configure the netboot connection by choosing one of the following actions.

**Note:** You should use the management port and IP as the netboot connection. Do not use a data LIF IP or a data outage might occur while the upgrade is being performed.

---

If Dynamic Host Configuration Protocol (DHCP) is...	Then...
Running	Configure the connection automatically by entering the following command at the boot environment prompt:  <b>ifconfig e0M -auto</b>

---



If Dynamic Host Configuration Protocol (DHCP) is...	Then...
Not running	<p>Manually configure the connection by entering the following command at the boot environment prompt:</p> <pre data-bbox="613 323 1414 380"><b>ifconfig e0M -addr=<i>filer_addr</i> mask=<i>netmask</i> -gw=<i>gateway</i> -dns=<i>dns_addr</i> domain=<i>dns_domain</i></b></pre> <p><i>filer_addr</i> is the IP address of the storage system.</p> <p><i>netmask</i> is the network mask of the storage system.</p> <p><i>gateway</i> is the gateway for the storage system.</p> <p><i>dns_addr</i> is the IP address of a name server on your network.</p> <p><i>dns_domain</i> is the Domain Name Service (DNS) domain name. If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL; you need only the server's host name.</p> <p><b>Note:</b> Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p>

19. Perform netboot by entering the following command on node3:

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

The *path\_to\_the\_web-accessible\_directory* should lead to where you downloaded the `netboot.tgz` in Step 1a in the section *Preparing for netboot*.

**Note:** Do not interrupt the boot.

20. From the boot menu, select option (7) `Install new software first`.

This menu option downloads and installs the new Data ONTAP image to the boot device.

**Note:** Disregard the following message: "This procedure is not supported for Non-Disruptive Upgrade on an HA pair". The note applies to nondisruptive upgrades of Data ONTAP, and not upgrades of controllers.

21. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL `http://path_to_the_web-accessible_directory/image.tgz`.

22. Complete the following substeps:

- a. Enter `n` to skip the backup recovery when you see the following prompt:

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

- b. Reboot by entering `y` when you see the following prompt:

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

The controller module reboots but stops at the boot menu because the boot device was reformatted and the configuration data needs to be restored.

23. Select (5) **Maintenance mode boot** by entering `5`, and then enter `y` when prompted to continue with the boot.

24. Before continuing to Step 26, go to the "Setting the FC or CNA configuration on node3" section to make any necessary changes to the FC or CNA ports on the node.

Make the changes recommended in those sections, reboot the node, and go into Maintenance mode.

25. Find the system ID of node3:

```
disk show -a
```

### Example

The system displays the system ID of the node and information about its disks, as shown in the following example:

```
*> disk show -a
Local System ID: 536881109
DISK      OWNER                                POOL  SERIAL NUMBER HOME                                DR HOME
-----
0b.02.23 nst-fas2520-2(536880939) Pool10 KPG2RK6F    nst-fas2520-2(536880939)
0b.02.13 nst-fas2520-2(536880939) Pool10 KPG3DE4F    nst-fas2520-2(536880939)
0b.01.13 nst-fas2520-2(536880939) Pool10 PPG4KLAA    nst-fas2520-2(536880939)
.....
0a.00.0   (536881109) Pool10 YFKSX6JG    (536881109)
.....
```

**Note:** You may see the message `disk show: No disks match option -a.` after entering the command. This is not an error message so you can continue with the procedure.

26. Reassign node1's spares, disks belonging to the root, and any non-root aggregates that were not relocated to node2 earlier in the "Relocating aggregates from node1 to node2" section:

Disk type ..	Run the command..
With shared disks	<code>disk reassign -s node1_sysid -d node3_sysid -p node2_sysid</code>
Without shared disks	<code>disk reassign -s node1_sysid -d node3_sysid</code>

For the `node1_sysid` value, use the information captured in Step 12 of the section "Recording node1 information". To obtain the value for `node3_sysid`, use the `sysconfig` command.

**Note:** The `-p` option is only required in maintenance mode when shared disks are present.

The `disk reassign` command reassigns only those disks for which `node1_sysid` is the current owner.

The system displays the following message:

```
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
```

27. Enter `n`.

The system displays the following message:

```
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
```

28. Enter `y`.

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
```

29. Enter **y**.

**Note:** After reassigning the disk run

**aggr status - r**  
command to verify that the root aggregate is in healthy state.

30. Verify that the controller and chassis are configured as **ha**:

**ha-config show**

### Example

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

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

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

If the controller and chassis are not configured as **ha**, use the following commands to correct the configuration: `ha-config modify controller ha` and `ha-config modify chassis ha`.

If you have a MetroCluster configuration, use the following commands to modify the controller and chassis: `ha-config modify controller mcc` and `ha-config modify chassis mcc`.

31. Destroy the mailboxes on node3:

**mailbox destroy local**

The console displays the following message:

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

32. Enter **y** at the prompt to confirm that you want to destroy the local mailboxes.

33. Exit Maintenance mode:

**halt**

The system stops at the boot environment prompt.

34. On node2, check the system date, time, and time zone:

**date**

35. On node3, check the date by entering the following command at the boot environment prompt:

**show date**

36. If necessary, set the date on node3:

**set date mm/dd/yyyy**

37. On node3, check the time by entering the following command at the boot environment prompt:

**show time**

38. If necessary, set the time on node3:

```
set time hh:mm:ss
```

39. Verify the partner system ID is set correctly as noted in step 26 on page 42 under -p switch:

```
printenv partner-sysid
```

40. If necessary, set the partner system ID on node3:

```
setenv partner-sysid node2_sysid
```

- a. Save the settings:

```
saveenv
```

41. Access the boot menu by entering the following command at the boot environment prompt:

```
boot_ontap menu
```

42. At the boot menu, select option (6) **Update flash from backup config** by entering 6 at the prompt.

The system displays the following message: This will replace all flash-based configuration with the last backup to disks. Are you sure you want to continue?:

43. Enter **y** at the prompt.

The boot proceeds normally, and the system then asks you to confirm the system ID mismatch.

**Note:** The system might reboot twice before displaying the mismatch warning.

44. Confirm the mismatch as shown in the following example:

#### Example

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

The node might go through one round of reboot before booting normally.

45. Log in to node3.

## Setting the FC or CNA configuration on node3

If node3 has onboard FC ports, onboard converged network adapter (CNA) ports, or a CNA card, you must configure the settings before completing the rest of the procedure.

### About this task

You might need to complete the section *Setting the onboard Fibre Channel configuration on node3*, the section *Configuring CNA ports on node3*, or both sections.

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

If node3 does not have onboard FC ports, onboard CNA ports, or a CNA card, and you are upgrading a system with storage disks, you can skip to the section *Mapping ports from node1 to node3*. However, if you have a V-Series system or a system with FlexArray Virtualization software with storage arrays, and node3 does not have onboard FC ports, onboard CNA ports, or a CNA card, return to the section *Installing and booting node3* and resume the section at Step 7.

### Choices

- [Configuring Fibre Channel ports on node3](#) on page 45
- [Checking and configuring CNA ports on node3](#) on page 47

## Configuring Fibre Channel ports on node3

If node3 has Fibre Channel (FC) ports—onboard or on an FC adapter—you must set port configurations on the node before you bring it into service because the ports are not preconfigured. If the ports are not configured, you might experience a disruption in service.

### Before you begin

You must have the values of the FC port settings from node1 that you saved in the section *Preparing the nodes for the upgrade*.

### About this task

You can skip this section if your system does not have FC configurations. If your system has onboard CNA ports or a CNA adapter, you configure them in the next section.

**Important:** If your system has storage disks, enter the commands in this section at the cluster prompt. If you have a V-Series system or have FlexArray Virtualization Software and are connected to storage arrays, enter commands in this section in Maintenance mode.

### Steps

1. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	Go to Step 5.
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Go to Step 2.

2. Boot node3 and access Maintenance mode by entering the following command:

```
boot_ontap maint
```

3. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	Enter the following command: <b>system node hardware unified-connect show</b>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Enter the following command: <b>ucadmin show</b>

The system displays information about all FC and converged network adapters on the system.

4. Compare the FC settings on node3 with the settings that you captured earlier from node1.

5. Take one of the following actions:

If the default FC settings on the new nodes are...	Then...
The same as the ones you that captured on node1	Go to Step 11.
Different from the ones that you captured on node1	Go to Step 6.

6. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	<p>Modify the FC ports on node3 as needed by entering one of the following commands:</p> <ul style="list-style-type: none"> <li>To program target ports: <code>system node hardware unified-connect modify -type   -t target -adapter port_name</code></li> <li>To program initiator ports: <code>system node hardware unified-connect modify -type   -t initiator -adapter port_name</code></li> </ul> <p>-type is the FC4 type: target or initiator.</p>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<p>Modify the FC ports on node3 as needed by entering the following command: <code>ucadmin modify -m fc -t initiator -f adapter_port_name</code></p> <p>-t is the FC4 type, target or initiator.</p> <p><b>Note:</b> The FC ports must be programmed as initiators.</p>

7. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	<p>Verify the new settings by entering the following command and examining the output:</p> <p><code>system node hardware unified-connect show</code></p>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<p>Verify the new settings by entering the following command and examining the output:</p> <p><code>ucadmin show</code></p>

8. Exit Maintenance mode by entering the following command:

`halt`

9. After you enter the command, wait until the system stops at the boot environment prompt.

10. Take one of the following actions:

If the system you are upgrading...	Then...
Is a V-Series system or has FlexArray Virtualization software running clustered Data ONTAP 8.3	<p>Boot node3 and access Maintenance mode by entering the following command at the boot environment prompt:</p> <p><code>boot_ontap maint</code></p>
Is not a V-Series system or does not have FlexArray Virtualization software	<p>Boot node3 by entering the following command at the boot environment prompt:</p> <p><code>boot_ontap</code></p>

11. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	<ul style="list-style-type: none"> <li>If node3 has a CNA card or CNA onboard ports, go to the section <i>Checking and configuring CNA ports on node3</i>.</li> <li>If node3 does not have a CNA card or CNA onboard ports, skip the section <i>Checking and configuring CNA ports node3</i> and go to the section <i>Mapping ports from node1 to node3</i>.</li> </ul>

If the system that you are upgrading...	Then...
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<ul style="list-style-type: none"> <li>If node3 has a CNA card or CNA onboard ports, go to the section <i>Checking and configuring CNA ports node3</i>.</li> <li>If node3 does not have a CNA card or CNA onboard ports, skip the section <i>Checking and configuring CNA ports node3</i>, return to the section <i>Installing and booting node3</i> and resume the section at Step 7.</li> </ul>

## Checking and configuring CNA ports on node3

If node3 has onboard CNA ports or a CNA card, you must check the configuration of the ports and possibly reconfigure them, depending on how you want to use the upgraded system.

### Before you begin

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

### About this task

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

CNA ports might be found on an adapter or on the controller, and have the following configurations, but you should check the configuration of the CNA ports on the node3 and change it, if necessary:

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

**Attention:** If your system has storage disks, you enter the commands in this section at the cluster prompt unless directed to enter Maintenance mode. If you have a V-Series system or have FlexArray Virtualization Software and are connected to storage arrays, you enter commands in this section at the Maintenance mode prompt. You must be in Maintenance mode to configure CNA ports.

### Steps

1. Check how the ports are currently configured by entering one of the following commands on node3:

If the system...	Then...
Has storage disks	<code>system node hardware unified-connect show</code>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<code>ucadmin show</code>

The system displays output similar to the following examples:

```
cluster1::> system node hardware unified-connect show
system node hardware unified-connect show
Node  Adapter  Current Current  Pending Pending  Admin
-----  -----  -----  -----  -----  -----  -----
f-a    0e         fc        initiator -         -         online
f-a    0f         fc        initiator -         -         online
f-a    0g         cna       target   -         -         online
f-a    0h         cna       target   -         -         online
f-b    0e         fc        initiator -         -         online
```

```
f-b 0f fc initiator - - online
f-b 0g cna target - - online
f-b 0h cna target - - online
12 entries were displayed.
```

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

2. If the current SFP+ module does not match the desired use, replace it with the correct SFP+ module.

Contact your NetApp representative to obtain the correct SFP+ module.

3. Examine the output of the `system node hardware unified-connect show` or `ucadmin show` command and determine whether the CNA ports have the personality you want.

4. Take one of the following actions:

If the CNA ports...	Then...
Do not have the personality that you want	Go to Step 5.
Have the personality that you want	Skip Step 5 through Step 12 and go Step 13.

5. Take one of the following actions:

If the system...	Then...
Has storage disks and is running clustered Data ONTAP 8.3	Boot node3 and enter Maintenance mode: <b>boot_ontap maint</b>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Go to Step 6. You should already be in Maintenance mode.

6. Take one of the following actions:

If you are configuring...	Then...
Ports on a CNA card	Go to Step 7.
Onboard CNA ports	Skip Step 7 and go to Step 8.

7. If the adapter is in initiator mode, and if the CNA port is online, take the CNA port offline:

```
storage disable adapter adapter_name
```

Adapters in target mode are automatically offline in Maintenance mode.

8. If the current configuration does not match the desired use, change the configuration as needed:

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

- `-m` is the personality mode, `fc` or `cna`.



- -t-t is the FC4 type, **target** or **initiator**.

**Note:** You need to use FC initiator for tape drives, FlexArray Virtualization systems, and MetroCluster configurations. You need to use the FC target for SAN clients.

9. Verify the settings:

```
ucadmin show
```

10. Take one of the following actions:

If the system...	Then...
Has storage disks	<p>a. Stop the system:</p> <pre>halt</pre> <p>The system stops at the boot environment prompt.</p> <p>b. Enter the following command:</p> <pre>boot_ontap</pre>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays and is running clustered Data ONTAP 8.3	<p>Reboot to Maintenance mode:</p> <pre>boot_ontap maint</pre>

11. Verify the settings:

If the system...	Then...
Has storage disks	<pre>system node hardware unified-connect show</pre>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<pre>ucadmin show</pre>

**Example**

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

```
cluster1::> system node hardware unified-connect show
Node  Adapter  Current Mode  Current Type  Pending Mode  Pending Type  Admin Status
----  -
f-a   1a         fc      initiator -      -      online
f-a   1b         fc      target  -      initiator online
f-a   2a         fc      target  cna    -      online
f-a   2b         fc      target  cna    -      online
4 entries were displayed.
```

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

12. Place any target ports online by entering one of the following commands, once for each port:

If the system...	Then...
Has storage disks	<code>network fcp adapter modify -node node_name -adapter adapter_name -state up</code>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<code>fcp config adapter_name up</code>

13. Cable the port.

14. Take one of the following actions:

If the system...	Then...
Has storage disks	Go to the “Mapping ports from node1 to node3” section.
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Return to the “Installing and booting node3” section and resume the section at Step 7.

### Mapping ports from node1 to node3

You need to make sure that the physical ports on node1 map correctly to the physical ports on node3, which will let node3 communicate with other nodes in the cluster and with the network after the upgrade.

#### Before you begin

You must already have information about the ports on the new nodes from the *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com). You use the information later in this section and in the “Mapping ports from node2 to node4” section.

The software configuration of node3 must match the physical connectivity of node3, and IP connectivity must be restored before you continue with the upgrade.

#### About this task

Port settings might vary, depending on the model of the nodes. You must make the original node's port and LIF configuration compatible with what you plan the new node's configuration to be. This is because the new node replays the same configuration when it boots, which means that when you boot node3, ONTAP will try to host LIFs on the same ports that were used on node1.

Therefore, if the physical ports on node1 do not map directly to the physical ports on node3, then software configuration changes will be required to restore cluster, management, and network connectivity after the boot. In addition, if the cluster ports on node1 do not directly map to the cluster ports on node3, node3 may not automatically rejoin quorum when it is rebooted until a software configuration change is made to host the cluster LIFs on the correct physical ports.

#### Steps

1. Record all the node1 cabling information for node1, the ports, broadcast domains, and IPspaces, in the following table:

LIF	Node1 ports	Node1 IPspaces	Node1 broadcast domains	Node3 ports	Node3 IPspaces	Node3 broadcast domains
Cluster 1						
Cluster 2						
Cluster 3						
Cluster 4						
Cluster 5						

LIF	Node1 ports	Node1 IPspaces	Node1 broadcast domains	Node3 ports	Node3 IPspaces	Node3 broadcast domains
Cluster 6						
Node management						
Cluster management						
Data 1						
Data 2						
Data 3						
Data 4						
SAN						
Intercluster port						

See the “Recording node1 information” section for the steps to obtain this information.

2. Record all the cabling information for node3, the ports, broadcast domains, and IPspaces in the previous table using the same procedure in “Recording node1 information” for the steps to obtain this information.
3. Follow these steps to verify if the setup is a two-node switchless cluster:
  - a. Set the privilege level to advanced:

#### Example

```
cluster::> set -privilege advanced

Warning: These advanced commands are potentially dangerous, use them only when
        directed to do so by NetApp personnel.
Do you wish to continue? (y or n): y

cluster::*>
```

- b. Verify if the setup is a two-node switchless cluster:

```
network options switchless-cluster show
```

#### Example

```
cluster::*> network options switchless-cluster show

Enable Switchless Cluster: false/true
```

The value of this command must match the physical state of the system.

- c. Return to the administration privilege level:

```
cluster::*> set -privilege admin

cluster::>
```

4. Get node3 into quorum by performing the following steps:
  - a. Boot node3. See the “Installing and booting node3” section to boot the node if you have not already done so.
  - b. Verify that the new cluster ports are in the Cluster broadcast domain:

```
network port show -node node -port port -fields broadcast-domain
```

#### Example

The following example shows that port e0a is in the Cluster domain on node3:

```
cluster::> network port show -node node3 -port e0a -fields broadcast-domain
node      port broadcast-domain
-----
node3     e1a Cluster
```

- c. Add the correct ports to the Cluster broadcast domain:

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

#### Example

This example adds Cluster port e1b on node3:

```
network port modify -node node3 -port e1b -ipSPACE Cluster -mtu 9000
```

**Note:** For a MetroCluster configuration, you may not be able to change the broadcast domain of a port because it is associated with a port hosting the LIF of a sync-destination SVM and see errors similar to, but not restricted to:

```
command failed: This operation is not permitted on a Vserver that is
                configured as the destination of a MetroCluster Vserver relationship.
```

Enter the following command from the corresponding sync-source SVM on the remote site to reallocate the sync-destination LIF to an appropriate port:

```
metrocluster vserver resync -vserver vserver_name
```

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

```
network interface migrate -vserver Cluster -lif lif_name -source-node node3 -destination-node node3 -destination-port port_name
```
  - e. Modify the home port of the cluster LIFs:
 

```
network interface modify -vserver Cluster -lif lif_name -home-port port_name
```
  - f. If the cluster ports are not in the Cluster broadcast-domain, add them with the following command:
 

```
network port broadcast-domain add-ports -ipSPACE Cluster -broadcast-domain Cluster -ports node:port
```
  - g. Remove the old ports from the Cluster broadcast domain:
 

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

#### Example

The following command removes port e0d on node3:

```
network port broadcast-domain remove-ports -ipSPACE Cluster -broadcast-domain Cluster -ports node3:e0d
```

- h. Verify that node3 has rejoined quorum:

```
cluster show -node node3 -fields health
```

- Adjust the broadcast domains hosting your cluster LIFs and node-management/cluster-management LIFs. Ensure that each broadcast domain contains the correct ports. A port cannot be moved between broadcast domains if it is hosting or is home to a LIF so you may need to migrate and modify the LIFs as follows:

- Display the home port of a LIF:

```
network interface show -fields home-node,home-port
```

- Display the broadcast domain containing this port:

```
network port broadcast-domain show -ports node_name:port_name
```

- Add or remove ports from broadcast domains:

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

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

- Modify a LIF's home port:

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

- Adjust the intercluster broadcast domains and migrate the intercluster LIFs, if necessary, using the same commands shown in Step 6.
- Adjust any other broadcast domains and migrate the data LIFs, if necessary, using the same commands shown in Step 5.
- If there were any ports on node1 that no longer exist on node3, follow these steps to delete them:

- Access the advanced privilege level on either node:

```
set -privilege advanced
```

The system displays the following message:

```
Warning: These advanced commands are potentially dangerous; use them only when
directed to do so by NetApp personnel.
do you wish to continue? (y or n):
```

- Enter **y**.

- To delete the ports:

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

- Return to the admin level:

```
set -privilege admin
```

- Adjust all the LIF failover groups:

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

### Example

The following command sets the failover policy to broadcast-domain-wide and uses the ports in failover group `fg1` as failover targets for LIF `data1` on node3:

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

See “Configuring failover settings on a LIF” in the *ONTAP 9 Network Management Guide* or *ONTAP 9 Commands: Manual Page Reference* for more information.

- Verify the changes on node3:

```
network port show -node node3
```

## Verifying the node3 installation

After you install and boot node3, you need to verify that it is installed correctly, and that node3 has joined the cluster and can communicate with node2.

### Steps

1. Verify that node3 is part of the same cluster as node2 and healthy by entering the following command:
2. Verify that node3 can communicate with node2 and that all LIFs are up by entering the following command:

```
cluster show
```

```
network interface show -curr-node node3
```

3. Take one of the following actions:

If the cluster is...	Then...
In a SAN environment	Complete Step 4 and then go to the section <i>Moving non-SAN data LIFs owned by node1 from node2 to node3 and verifying SAN LIFs on node3.</i>
Not in a SAN environment	Skip Step 4 and go to the section <i>Moving non-SAN data LIFs owned by node1 from node2 to node3 and verifying SAN LIFs on node3.</i>

4. Verify that node2 and node3 are in quorum by entering the following command on one of the nodes and examining its output:

```
event log show -messagename scsiblade.*
```

### Example

The following example shows the output when the nodes in the cluster are in quorum:

```
cluster::> event log show -messagename scsiblade.*
Time           Node           Severity      Event
-----
8/13/2012 14:03:51 node1          INFORMATIONAL scsiblade.in.quorum: The scsi-blade...
8/13/2012 14:03:51 node2          INFORMATIONAL scsiblade.in.quorum: The scsi-blade...
8/13/2012 14:03:48 node3          INFORMATIONAL scsiblade.in.quorum: The scsi-blade...
8/13/2012 14:03:43 node4          INFORMATIONAL scsiblade.in.quorum: The scsi-blade...
```

## Moving non-SAN data LIFs owned by node1 from node2 to node3 and verifying SAN LIFs on node3

After you verify the node3 installation and before you relocate aggregates from node2 to node3, you need to move the non-SAN data LIFs belonging to node1 that are currently on node2 from node2 to node3. You also need to verify the SAN LIFs on node3.

### About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. Moving SAN LIFs is not necessary for cluster or service health during the upgrade. SAN LIFs are not moved unless they need to be mapped to new ports. You will verify that the LIFs are healthy and located on appropriate ports after you bring node3 online.

### Steps

1. List all the non-SAN data LIFs not owned by node2 by entering the following command on either node and capturing the output:

```
network interface show -role data -curr-node node2 -is-home false -home-node node3
```

2. Take one of the following actions:

If node1...	Description
Had interface groups or VLANs configured	Go to Step 3.
Did not have interface groups or VLANs configured	Skip Step 3 and go to Step 4.

3. Perform the following substeps to migrate any non-SAN data LIFs hosted on interface groups and VLANs that were originally on node1 from node2 to node3:

- a. Migrate any data LIFs hosted on node2 that previously belonged to node1 on an interface group to a port on node3 that is capable of hosting LIFs on the same network by entering the following command, once for each LIF:

```
network interface migrate -vserver vservice_name -lif lif_name -destination-node node3 -
destination-port netport|ifgrp
```

- b. Modify the home port and home node of the LIF in substep a to the port and node currently hosting the LIFs by entering the following command, once for each LIF:

```
network interface modify -vserver vservice_name -lif lif_name -home-node node3 -home-port
netport|ifgrp
```

- c. Migrate any data LIF hosted on node2 that previously belonged to node1 on a VLAN port to a port on node3 that is capable of hosting LIFs on the same network by entering the following command, once for each LIF:

```
network interface migrate -vserver vservice_name -lif lif_name -destination-node node3 -
destination-port netport|ifgrp
```

- d. Modify the home port and home node of the LIFs in substep c to the port and node currently hosting the LIFs by entering the following command, once for each LIF:

```
network interface modify -vserver vservice_name -lif lif_name -home-node node3 -home-port
netport|ifgrp
```

4. Take one of the following actions:

If the cluster is configured for...	Then...
NAS	Complete Step 5 and Step 6, skip Step 7, and complete Step 8 through Step 11.
SAN	Disable all the SAN LIFs on the node to take them down for the upgrade by entering the following command:  <pre>network interface modify -vserver vservice_name -lif lif_name - home-node node_to_upgrade -home-port netport ifgrp -status- admin down</pre> Repeat this command on the other node in the HA pair.

5. If you have data ports that are not the same on your platforms, enter the following command to add the ports to the broadcast domain:

```
network port broadcast-domain add-ports -ip-space IPspace_name -broadcast-domain mgmt -ports
node:port
```

#### Example

The following example adds port e0a on node 6280-1 and port e0i on node 8060-1 to broadcast domain mgmt in the IPspace Default:

```
cluster::> network port broadcast-domain add-ports -ip-space Default
-broadcast-domain mgmt -ports 6280-1:e0a, 8060-1:e0i
```

6. Migrate each non-SAN data LIF to node3 by entering the following command, once for each LIF:

```
network interface migrate -vserver vserver_name -lif lif_name -destination-node node3 -
destination-port netport|ifgrp
```

7. Make sure that the data migration is persistent by entering the following command:

```
network interface modify -vserver vserver_name -lif lif_name -home-port netport|ifgrp -home-
node node3
```

8. Ensure that the SAN LIFs are on the correct ports on node3 by completing the following substeps:

- a. Enter the following command and examine its output:

```
network interface show -data-protocol iscsi|fcip -home-node node3
```

### Example

The system returns output similar to the following example:

```
cluster::> net int show -data-protocol iscsi|fcip -home-node node3
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
vs0	a0a	up/down	10.63.0.53/24	node3	a0a	true
	data1	up/up	10.63.0.50/18	node3	e0c	true
	rads1	up/up	10.63.0.51/18	node3	e1a	true
	rads2	up/down	10.63.0.52/24	node3	e1b	true
vs1	lif1	up/up	172.17.176.120/24	node3	e0c	true
	lif2	up/up	172.17.176.121/24	node3	e1a	true

- b. If node3 has any SAN LIFs or groups of SAN LIFs that are on a port that did not exist on node1 or that need to be mapped to a different port, move them to an appropriate port on node3 by completing the following substeps:

- i. Set the LIF status to down by entering the following command:

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

- ii. Remove the LIF from the port set:

```
portset remove -vserver vserver_name -portset portset_name -port-name port_name
```

- iii. Enter one of the following commands:

- Move a single LIF by entering the following command:

```
network interface modify -vserver vserver_name -lif lif_name -home-port
new_home_port
```

- Move all the LIFs on a single nonexistent or incorrect port to a new port by entering the following command:

```
network interface modify {-home-port port_on_node1 -home-node node1 -role data} -
home-port new_home_port_on_node3
```

- Add the LIFs back to the port set:

```
portset add -vserver vserver_name -portset portset_name -port-name port_name
```

**Note:** You need to ensure that you move SAN LIFs to a port that has the same link speed as the original port.

9. Modify the status of all LIFs to “up” so the LIFs can accept and send traffic on the node by entering the following command:

```
network interface modify -home-port port_name -home-node node3 -lif data -status-admin up
```

10. Enter the following command and examine its output to verify that LIFs have been moved to the correct ports and that the LIFs have the status of “up” by entering the following command on either node and examining the output:

```
network interface show -home-node node3 -role data
```



11. 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 -status-admin up`
12. Send a post-upgrade AutoSupport message to NetApp for node1 by entering the following command:  
`system node autosupport invoke -node node3 -type all -message "node1 successfully upgraded from platform_old to platform_new"`

## Relocating non-root aggregates from node2 to node3

Before you can replace node2 with node4, you must send an AutoSupport message for node2 and then relocate the non-root aggregates that are owned by node2 to node3.

### Steps

1. Send an AutoSupport message to NetApp for node2:

```
system node autosupport invoke -node node2 -type all -message "Upgrading node2 from platform_old to platform_new"
```

2. Verify that the AutoSupport message was sent:

```
system node autosupport show -node node2 -instance
```

The fields `Last Subject Sent:` and `Last Time Sent:` contain the message title of the last message that was sent and the time when the message was sent.

3. Relocate the non-root aggregates:

- a. Set the privilege level to advanced:

```
set -privilege advanced
```

The system displays the following message:

```
Warning: These advanced commands are potentially dangerous; use them only when
directed to do so by NetApp personnel.
do you wish to continue? {y or n}:
```

- b. Enter `y`.
- c. List the aggregates that are owned by node2:

```
storage aggregate show -owner-name node2
```

- d. Start aggregate relocation:

```
storage aggregate relocation start -node node2 -destination node3 -aggregate-list * -ndo-controller-upgrade true
```

**Note:** The command locates only non-root aggregates.

The system displays the following message:

```
Warning: Aggregate relocation will not change home ownership of an aggregate
which is owned by the source node during an NDO controller upgrade. This
parameter should be used only while performing NDO controller upgrade.
Are you performing NDO controller upgrade? {y|n}:
```

- e. When prompted, enter `y`.

Relocation occurs in the background. It can take anywhere from a few seconds to a couple of minutes to relocate an aggregate. The time includes both client outage and non-outage portions. The command does not relocate any offline or restricted aggregates.

f. Return to the admin privilege level:

```
set -privilege admin
```

4. Verify the relocation status of node2:

```
storage aggregate relocation show -node node2
```

The output will display Done for an aggregate after it has been relocated.

**Note:** You must wait until all of the aggregates that are owned by node2 have been relocated to node3 before proceeding to the next step.

5. Take one of the following actions:

If relocation of...	Then...
All aggregates was successful	Go to Step 6.
Any aggregates failed, or was vetoed	<p>a. Check the EMS logs for the corrective action.</p> <p>b. Perform the corrective action.</p> <p>c. Set the privilege level to advanced: <b>set -privilege advanced</b> The system displays the following message:</p> <pre>Warning: These advanced commands are potentially dangerous; use them only when directed to do so by NetApp personnel. do you wish to continue? {y or n}:</pre> <p>d. Enter <b>y</b>.</p> <p>e. Relocate any failed or vetoed aggregates: <b>storage aggregate relocation start -node node2 -destination node3 -aggregate-list * -ndo-controller-upgrade true</b> The system displays the following message:</p> <pre>Warning: Aggregate relocation will not change home ownership of an aggregate which is owned by the source node during an NDO controller upgrade. This parameter should be used only while performing NDO controller upgrade. Are you performing NDO controller upgrade? {y n}:</pre> <p>f. When prompted, enter <b>y</b>.</p> <p>g. Return to the admin privilege level: <b>set -privilege admin</b></p>

If necessary, you can force the relocation by using one of the following methods:

- By overriding veto checks using the `storage aggregate relocation start -override-vetoes true -ndo-controller-upgrade` command
- By overriding destination checks using the `storage aggregate relocation start -override-destination-checks true -ndo-controller-upgrade` command

See the *ONTAP 9 Disks and Aggregates Power Guide* and the *ONTAP 9 Commands: Manual Page Reference* for more information about the storage aggregate relocation commands.

- Verify that all of the non-root aggregates are online on node3:

```
storage aggregate show -node node3 -state offline -root -false
```

If any aggregates have gone offline or have become foreign, you must bring them online, once for each aggregate:

```
storage aggregate online -aggregate aggr_name
```

- Verify that all of the volumes are online on node3:

```
volume show -node node3 -state offline
```

If any volumes are offline on node3, you must bring them online, once for each volume:

```
volume online -vserver vservice-name -volume volume-name
```

- Verify that node2 does not own any online non-root aggregates:

```
storage aggregate show -owner-name node2 -ha-policy sfo -state online
```

The command output should not display online non-root aggregates because all of the non-root online aggregates have already been relocated to node3.

### Moving non-SAN data LIFs owned by node2 to node3

After you relocate the aggregates from node2 to node3, you need to move the non-SAN data LIFs owned by node2 to node3. The method you use depends on whether the cluster is configured for NAS or SAN.

#### About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. Moving SAN LIFs is not necessary for cluster or service health during the upgrade. SAN LIFs are not moved unless they need to be mapped to new ports. You must verify that the LIFs are healthy and located on the appropriate ports after you move the LIFs from node3 to node4 and bring node4 online.

#### Steps

- List all the non-SAN data LIFs owned by node2 by entering the following command on either node and capturing the output:

```
network interface show -data-protocol nfs|cifs -home-node node2
```

#### Example

The following example shows the command output for node2:

```
cluster::> network interface show -data-protocol nfs|cifs -home-node node2
Logical      Status      Network      Current      Current      Is
Vserver      Interface   Admin/Oper   Address/Mask Node          Port         Home
-----
vs0
  a0a         up/down    10.63.0.53/24 node2        a0a          true
  data1       up/up      10.63.0.50/18 node2        e0c          true
  rads1       up/up      10.63.0.51/18 node2        e1a          true
  rads2       up/down    10.63.0.52/24 node2        e1b          true
vs1
  lif1        up/up      172.17.176.120/24 node2        e0c          true
  lif2        up/up      172.17.176.121/24 node2        e1a          true
```

- Take one of the following actions:

If node2...	Description
Has interface groups or VLANs configured	Go to Step 3.
Does not have interface groups or VLANs configured	Skip Step 3 and go to Step 4.

3. Take the following steps to migrate non-SAN data LIFs hosted on interface groups and VLANs on node2:

- a. Migrate any data LIFs hosted on an interface group on node2 to a port on node3 that is capable of hosting LIFs on the same network by entering the following command, once for each LIF:

```
network interface migrate -vserver vservice_name -lif lif_name -destination-node node3 -
destination-port netport|ifgrp
```

- b. Modify the home port and home node of the LIFs in substep a to the port and node currently hosting the LIFs by entering the following command, once for each node:

```
network interface modify -vserver vservice_name -lif lif_name -home-node node3 -home-port
netport|ifgrp
```

- c. Migrate any LIFs hosted on VLANs on node2 to a port on node3 that is capable of hosting LIFs on the same network as that of the VLANs by entering the following command, once for each LIF:

```
network interface migrate -vserver vservice_name -lif lif_name -destination-node node3 -
destination-port netport|ifgrp
```

- d. Modify the home port and home node of the LIFs in substep c to the port and node currently hosting the LIFs by entering the following command, once for each LIF:

```
network interface modify -vserver vservice_name -lif lif_name -home-node node3 -home-port
netport|ifgrp
```

4. Take one of the following actions:

If the cluster is configured for...	Then...
NAS	Complete Step 5 through Step 8.
SAN	Skip Step 5 through Step 8 and then complete Step 9.
Both NAS and SAN	Complete Step 5 through Step 9.

5. If you have data ports that are not the same on your platforms, enter the following command to add the ports to the broadcast domain:

```
network port broadcast-domain add-ports -ip-space IPspace_name -broadcast-domain mgmt -ports
node:port
```

#### Example

The following example adds port e0a on node 6280-1 and port e0i on node 8060-1 to broadcast domain mgmt in the IPspace Default:

```
cluster::> network port broadcast-domain add-ports -ip-space Default
-broadcast-domain mgmt -ports 6280-1:e0a, 8060-1:e0i
```

6. Migrate each non-SAN data LIF to node3 by entering the following command, once for each LIF:

```
network interface migrate -vserver vservice_name -lif lif_name -destination-node node3 -
destination-port netport|ifgrp
```

7. Verify that non-SAN LIFs have been moved to the correct ports and that the LIFs have the status of up by entering the following command on either node and examining the output:

```
network interface show -curr-node node3 -data-protocol cifs|nfs
```

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 vservice_name -lif lif_name -status-admin up
```

9. If you have interface groups or VLANs configured, complete the following substeps:

- a. Remove the VLANs from the interface groups by entering the following command:

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

- b. Enter the following command and examine its output to see if there are any interface groups configured on the node:

```
network port ifgrp show -node node_name -ifgrp ifgrp_name -instance
```

The system displays interface group information for the node, as shown in the following example:

```
cluster::> network port ifgrp show -node node2 -ifgrp a0a -instance
      Node: node2
Interface Group Name: a0a
Distribution Function: ip
      Create Policy: multimode_lacp
      MAC Address: 02:a0:98:17:dc:d4
Port Participation: partial
      Network Ports: e2c, e2d
      Up Ports: e2c
      Down Ports: e2d
```

- c. If any interface groups are configured on the node, record the names of the interface groups and the ports assigned to them and then delete the ports by entering the following command, once for each port:

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

## Upgrading the node pair: Stage 4

During Stage 4, you record node2 information and then retire node2.

### Steps

1. [Recording node2 information](#) on page 61
2. [Retiring node2](#) on page 63

### Recording node2 information

Before you can shut down and retire node2, you need to record information about its cluster network, management, and FC ports as well as its NVRAM System ID. You need that information later in the procedure when you map node2 to node4 and reassign disks.

### Steps

1. Find the cluster network, node-management, intercluster, and cluster-management ports on node2:

```
network interface show -curr-node node_name -role cluster,intercluster,node-mgmt,cluster-mgmt
```

### Example

The system displays the LIFs for that node and other nodes in the cluster, as shown in the following example:

```
cluster::> network interface show -curr-node node2 -role cluster,intercluster,node-
mgmt,cluster-mgmt
      Logical      Status      Network      Current      Current      Is
Vserver  Interface  Admin/Oper  Address/Mask  Node         Port         Home
-----
node2
intercluster up/up      192.168.1.202/24  node2        e0e          true
clus1      up/up      169.254.xx.xx/24  node2        e0a          true
clus2      up/up      169.254.xx.xx/24  node2        e0b          true
mgmt1      up/up      192.168.0.xxx/24  node2        e0c          true
4 entries were displayed.
```

**Note:** Your system might not have intercluster LIFs. You will have a cluster management LIF only on one node of a node pair. A cluster management LIF was displayed in the example output of Step 1 in the section *Recording node1 port information*.

2. Capture the information in the output to use in the section *Mapping ports from node2 to node4*.

The output information is required to map the new controller ports to the old controller ports.

3. Determine physical ports on node2:

```
network port show -node node_name -type physical
```

*node\_name* is the node which is being migrated.

### Example

The system displays the physical ports on node2, as shown in the following example:

```
cluster::> network port show -node node2 -type physical
```

Node	Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Oper
node2	e0M	Default	172.17.178.19/24	up	1500	auto/100
	e0a	Default	-	up	1500	auto/1000
	e0b	Default	-	up	1500	auto/1000
	e1a	Cluster	Cluster	up	9000	auto/10000
	e1b	Cluster	Cluster	up	9000	auto/10000

5 entries were displayed.

4. Record the ports and their broadcast domains.

The broadcast domains will need to be mapped to the ports on the new controller later in the procedure.

5. Determine the FC ports on node2:

```
network fcp adapter show
```

### Example

The system displays the FC ports on the node2, as shown in the following example:

```
cluster::> network fcp adapter show -node node2
```

Node	Adapter	Connection Established	Host Port Address
node2	0a	ptp	11400
node2	0c	ptp	11700
node2	6a	loop	0
node2	6b	loop	0

4 entries were displayed.

6. Record the ports.

The output information is required to map the new FC ports on the new controller later in the procedure.

7. If you have not done so earlier, check whether there are interface groups or VLANs configured on node2:

```
ifgrp show
```

```
vlan show
```

You will use the information in the section *Mapping ports from node2 to node4*.

8. Take one of the following actions:

If you...	Then...
Recorded NVRAM System ID number in the section <i>Preparing the nodes for the upgrade</i>	Go to the next section, <i>Retiring node2</i> .
Did not record the NVRAM System ID number in the section <i>Preparing the nodes for the upgrade</i>	Complete Step 9 and Step 10 and then go to the next section, <i>Retiring node2</i> .

9. To display the attributes of node 2, use the

**system node show -instance -node node2**  
command.

```
cluster::> system node show -instance -node node2
...
NVRAM System ID: 1873808129
...
```

10. Record the NVRAM System ID to use in the section *Installing and booting node4*.

## Retiring node2

To retire node2, you need to shut node2 down properly and remove it from the rack or chassis. If the cluster is in a SAN environment, you also need to delete the SAN LIFs.

### Steps

1. Take one of the following actions:

If the cluster is...	Description
A two-node cluster	Go to Step 2.
A cluster with more than two nodes	Go to Step 9.

2. Access the advanced privilege level by entering the following command on either node:

**set -privilege advanced**

The system displays the following message:

```
Warning: These advanced commands are potentially dangerous; use them only when directed
to do so by NetApp personnel.
do you wish to continue? (y or n):
```

3. Enter **y**.
4. Verify that the cluster HA has been disabled by entering the following command and examining its output:

**cluster ha show**

The system displays the following message:

```
High Availability Configured: false
```

5. Check if node2 currently holds epsilon by entering the following command:

**cluster show**

The following example shows that node2 holds epsilon :

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

Warning: Cluster HA has not been configured. Cluster HA must be configured on a
two-node cluster to ensure data access availability in the event of
storage failover. Use the "cluster ha modify -configured true" command
to configure cluster HA.

2 entries were displayed.
```

6. If node2 holds epsilon, mark epsilon as **false** on the node so that it can be transferred to node3 by entering the following command:

```
cluster modify -node node2 -epsilon false
```

7. Transfer epsilon to node3 by marking epsilon **true** on node3 by entering the following command:

```
cluster modify -node node3 -epsilon true
```

8. Verify if the setup is a two-node switchless cluster:

```
network options switchless-cluster show
```

#### Example

```
cluster*:*> network options switchless-cluster show
Enable Switchless Cluster: false/true
```

The value of this command must match the physical state of the system.

9. Verify if the setup is a two-node switchless cluster:

```
network options switchless-cluster show
```

#### Example

```
cluster*:*> network options switchless-cluster show
Enable Switchless Cluster: false/true
```

The value of this command must match the physical state of the system.

10. Return to the admin level by entering the following command:

```
set -privilege admin
```

11. Halt node2 by entering the following command on either controller:

```
system node halt -node node2
```

12. After node2 shuts down completely, remove it from the chassis or the rack.

You can decommission node2 after the upgrade is completed. See *Decommissioning the old system*.



## Upgrading the node pair: Stage 5

During Stage 5, you install and boot node4, map the cluster and node-management ports from node2 to node4, and verify the node4 installation. You also move the data LIFs and SAN LIFs owned by node2 from node3 to node4, and relocate node2's aggregates from node3 to node4.

### Steps

1. [Installing and booting node4](#) on page 65
2. [Setting the FC or CNA configuration on node4](#) on page 71
3. [Mapping ports from node2 to node4](#) on page 77
4. [Verifying the node4 installation](#) on page 81
5. [Moving non-SAN data LIFs owned by node2 from node3 to node4 and verifying SAN LIFs on node4](#) on page 82
6. [Relocating node2's non-root aggregates from node3 to node4](#) on page 85

### Installing and booting node4

You need to install node4 in the rack, transfer node2 connections to node4, and boot node4. You must also reassign any node2 spares, any disks belonging to root, and any non-root aggregates that were not relocated to node3 earlier.

#### About this task

You need to netboot node4 if it does not have the same or later version of Data ONTAP 8.3 installed on it that is installed on node2. After you install node4, you can boot it from the Data ONTAP 8.3 image stored on the web server. You can then download the correct files to the boot media device for subsequent system boots.

However, you do not need to netboot node4 if it has the same or later version of Data ONTAP 8.3 that is installed on node2.

**Important:** If you are upgrading a V-Series system or a system with FlexArray Virtualization Software that is connected to storage arrays, you need to complete Step 1 through Step 7, leave this section at Step 8 and follow instructions in the section *Setting the FC or CNA configuration on node4* as needed, entering the commands in Maintenance mode. You then need to return to this section and resume the procedure with Step 9.

However, if you are upgrading a system with storage disks, you need to complete this entire section and then proceed to the section *Setting the FC or CNA configuration on node4*, entering commands at the cluster prompt.

### Steps

1. Take one of the following actions:

If node4 will be in ...	Then...
A chassis separate from node3	Go to Step 2.
The same chassis with node3	Skip Step 2 and Step 3 and go to Step 4.

2. Make sure that node4 has sufficient rack space.

If node4 is in a separate chassis from node3, you can put node4 in the same location as node2. If node3 and node4 are in the same chassis, then node4 is already in its appropriate rack location.

3. Install node4 in the rack, following the instructions in the *Installation and Setup Instructions* for the node model.
4. Cable node4, moving the connections from node2 to node4.

For systems with disks, cable the following connections using the instructions in the *Installation and Setup Instructions* or the *FlexArray Virtualization Installation Requirements and Reference* for the node4 platform, the appropriate disk shelf guide, and the *ONTAP 9 High-Availability Configuration Guide*:

- Console (remote-management port)
- Cluster ports
- Data ports
- Cluster and node-management ports
- Storage
- SAN configurations: iSCSI Ethernet and FC switch ports

**Note:** You do not need to move the interconnect card/FC\_VI card or interconnect/FC\_VI cable connection from node2 to node4 because most platform models have unique interconnect card models.

5. Take one of the following actions:

If node4 is in...	Then...
The same chassis as node3	Go to Step 8.
A chassis separate from node3	Go to Step 6.

6. Turn on the power to node4, and then interrupt the boot by pressing Ctrl-C to access the boot environment prompt.

**Attention:** When you boot node4, you might see the following message:

```
WARNING: The battery is unfit to retain data during a power
         outage. This is likely because the battery is
         discharged but could be due to other temporary
         conditions.
         When the battery is ready, the boot process will
         complete and services will be engaged.
         To override this delay, press 'c' followed by 'Enter'
```

7. If you see the warning message in Step 6, take the following actions:

- Check for any console messages that might indicate a problem other than a low NVRAM battery and, if necessary, take any required corrective action.
- Allow the battery to charge and the boot process to finish.

**Attention:** Do not override the delay. Failure to allow the battery to charge could result in a loss of data.

8. Take one of the following actions:

If your system...	Description
Has disks and no back-end storage	Skip Step 9 through Step 14 and go to Step 15.
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<ol style="list-style-type: none"> <li>Go to the section <i>Setting the FC or CNA configuration on node4</i> and complete the subsections in it, <i>Configuring Fibre Channel ports on node4</i> and <i>Checking and configuring CNA ports on node4</i>, as appropriate to your system.</li> <li>Return to this section and complete the remaining steps, beginning with Step 9.</li> </ol> <p><b>Important:</b> You must reconfigure FC onboard ports, CNA onboard ports, and CNA cards before you boot Data ONTAP on the V-Series system.</p>

9. Add the FC initiator ports of the new node to the switch zones.

See your storage array and zoning documentation for instructions.

10. Add the FC initiator ports to the storage array as new hosts, mapping the array LUNs to the new hosts.

See your storage array and zoning documentation for instructions.

11. Modify the World Wide Port Name (WWPN) values in the host or volume groups associated with array LUNs on the storage array.

Installing a new controller module changes the WWPN and World Wide Node Name (WWNN) values associated with each onboard FC port.

12. If your configuration uses switch-based zoning, adjust the zoning to reflect the new WWNN values.
13. Verify that the array LUNs are now visible to node4 by entering the following command and examining its output:

```
sysconfig -v
```

The system displays all of the array LUNs visible to each of the FC initiator ports. If the array LUNs are not visible, you cannot reassign disks from node2 to node4 later in this section.

14. At the Maintenance mode prompt, enter the following command:

```
halt
```

The system stops at the boot environment prompt.

15. Configure node4 for ONTAP by entering the following command:

```
set-defaults
```

16. If FDE is in use on this configuration, the `setenv bootarg.storageencryption.supportvariable` must be set to **true**, and the `kmip.init.maxwait` variable needs to be set to **off** to avoid a boot loop after the node2 configuration is loaded:

```
setenv bootarg.storageencryption.support true
```

```
setenv kmip.init.maxwait off
```

17. If the version of Data ONTAP installed on node4 is the same or later than the version of Data ONTAP 8.3 installed on node2, enter the following command:

```
boot_ontap menu
```

18. Take one of the following actions:

If the system you are upgrading...	Then...
Does <i>not</i> have the correct or current Data ONTAP version on node3	Go to Step 19.
Has the correct or current version of Data ONTAP on node3	Go to Step 24.

19. Configure the netboot connection by choosing one of the following actions.

**Note:** You should use the management port and IP as the netboot connection. Do not use a data LIF IP or a data outage might occur while the upgrade is being performed.

If Dynamic Host Configuration Protocol (DHCP) is...	Then...
Running	Configure the connection automatically by entering the following command at the boot environment prompt: <pre>ifconfig e0M -auto</pre>

If Dynamic Host Configuration Protocol (DHCP) is...	Then...
Not running	<p>Manually configure the connection by entering the following command at the boot environment prompt:</p> <pre data-bbox="613 323 1414 380"><b>ifconfig e0M -addr=<i>filer_addr</i> mask=<i>netmask</i> -gw=<i>gateway</i> -dns=<i>dns_addr</i> domain=<i>dns_domain</i></b></pre> <p><i>filer_addr</i> is the IP address of the storage system.</p> <p><i>netmask</i> is the network mask of the storage system.</p> <p><i>gateway</i> is the gateway for the storage system.</p> <p><i>dns_addr</i> is the IP address of a name server on your network.</p> <p><i>dns_domain</i> is the Domain Name Service (DNS) domain name. If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL; you need only the server's host name.</p> <p><b>Note:</b> Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p>

20. Perform netboot by entering the following command on node4:

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

The *path\_to\_the\_web-accessible\_directory* should lead to where you downloaded the `netboot.tgz` in Step 1a in the section *Preparing for netboot*.

**Note:** Do not interrupt the boot.

21. From the boot menu, select option (7) **Install new software first**.

This menu option downloads and installs the new Data ONTAP image to the boot device.

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

22. If you are prompted to continue the procedure, enter **y**, and when prompted for the package, enter the URL `http://path_to_the_web-accessible_directory/image.tgz`.
23. Complete the following substeps:

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

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

- b. Reboot by entering **y** when you see the following prompt:

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

The controller module reboots but stops at the boot menu because the boot device was reformatted and the configuration data needs to be restored.

24. Select (5) **Maintenance mode boot** by entering 5, and then enter **y** when prompted to continue with the boot.
25. Before continuing to Step 26, go to the section "Setting the FC or CNA configuration on node3" to make any necessary changes to the FC or CNA ports on the node.

Make the changes recommended in those sections, reboot the node, and go into Maintenance mode.

26. Enter the following command and examine the output to find the system ID of node4:

```
disk show -a
```

### Example

The system displays the system ID of the node and information about its disks, as shown in the following example:

```
*> disk show -a
Local System ID: 536881109
DISK      OWNER
-----
0b.02.23  nst-fas2520-2(536880939)  Pool0  KPG2RK6F  nst-fas2520-2(536880939)
0b.02.13  nst-fas2520-2(536880939)  Pool0  KPG3DE4F  nst-fas2520-2(536880939)
0b.01.13  nst-fas2520-2(536880939)  Pool0  PPG4KLAA  nst-fas2520-2(536880939)
.....
0a.00.0   (536881109)              Pool0  YFKSX6JG  (536881109)
.....
```

27. Reassign node2's spares, disks belonging to the root, and any non-root aggregates that were not relocated to node3 earlier in the "Relocating aggregates from node2 to node3" section:

Disk type ..	Run the command..
With shared disks	<b>disk reassign -s</b> <i>node2_sysid -d node4_sysid -p node3_sysid</i>
Without shared disks	<b>disk reassign -s</b> <i>node2_sysid -d node3_sysid</i>

For the *node2\_sysid* value, use the information captured in Step 10 of the "Recording node2 port information" section. For *node4\_sysid*, use the information captured in Step 23.

**Note:** The *-p* option is only required in maintenance mode when shared disks are present.

The `disk reassign` command will reassign only those disks for which *node2\_sysid* is the current owner.

The system displays the following message:

```
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
```

28. Enter **n**.

The system displays the following message:

```
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
```

29. Enter **y**.

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
```

30. Enter **y**.

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

**ha-config show**

### Example

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

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

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

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

If you have a MetroCluster configuration, use the following commands to correct the configuration: `ha-config modify controller mcc` and `ha-config modify chassis mcc`.

32. Destroy the mailboxes on node4:

```
mailbox destroy local
```

The console displays the following message:

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

33. Enter `y` at the prompt to confirm that you want to destroy the local mailboxes.

34. Exit Maintenance mode:

```
halt
```

The system stops at the boot environment prompt.

35. On node3, check the system date, time, and time zone:

```
date
```

36. On node4, check the date at the boot environment prompt:

```
show date
```

37. If necessary, set the date on node4:

```
set date mm/dd/yyyy
```

38. On node4, check the time at the boot environment prompt:

```
show time
```

39. If necessary, set the time on node4:

```
set time hh:mm:ss
```

40. Verify the partner system ID is set correctly as noted in step 27 on page 69 under `-p` switch.

```
printenv partner-sysid
```

41. If necessary, set the partner system ID on node4:

```
setenv partner-sysid node3_sysid
```

- a. Save the settings:

```
saveenv
```

42. Enter the boot menu at the boot environment prompt:

```
boot_ontap menu
```

43. At the boot menu, select option **(6) Update flash from backup config** by entering **6** at the prompt.

The system displays the following message: This will replace all flash-based configuration with the last backup to disks. Are you sure you want to continue?:

44. Enter **y** at the prompt.

The boot proceeds normally, and the system prompts you to confirm the system ID mismatch.

**Note:** The system might reboot twice before displaying the mismatch warning.

45. Confirm the mismatch, as shown in the following example:

#### Example

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

The node might complete one round of rebooting before booting normally.

46. Log in to node4.

## Setting the FC or CNA configuration on node4

If node4 has onboard FC ports, onboard converged network adapter (CNA) ports, or a CNA card, you must configure the settings before completing the rest of the procedure.

### About this task

You might need to complete the “Setting the onboard Fibre Channel configuration on node4” section, the “Configuring CNA ports on node4” section, or both sections.

If node4 does not have onboard FC ports, onboard CNA ports, or a CNA card, and you are upgrading a system with storage disks, you can skip to the “Mapping ports from node2 to node4” section. However, if you have a V-Series system or have FlexArray Virtualization Software and are connected to storage arrays, and node4 does not have onboard FC ports, onboard CNA ports, or a CNA card, you must return to the “Installing and booting node4” section and resume the section at Step 9.

### Choices

- [Configuring Fibre Channel ports on node4](#) on page 71
- [Checking and configuring CNA ports on node4](#) on page 74

## Configuring Fibre Channel ports on node4

If node4 has Fibre Channel (FC) ports onboard or on an FC adapter, you must set port configurations on the node before you bring it into service because the ports are not preconfigured. If the ports are not configured, you might experience a disruption in service.

### Before you begin

You must have the values of the FC port settings from node2 that you saved in the section *Preparing the nodes for the upgrade*.

## About this task

You can skip this section if your system does not have FC configurations. If your system has onboard CNA ports or a CNA adapter, you configure them in the next section.

**Important:** If your system has storage disks, you must enter the commands in this section at the cluster prompt. If you have a V-Series system or a system with FlexArray Virtualization Software connected to storage arrays, you enter commands in this section in Maintenance mode.

## Steps

1. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	Go to Step 5.
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Go to Step 2.

2. Access Maintenance mode by entering the following command:

```
boot_ontap maint
```

3. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	Enter the following command: <b>system node hardware unified-connect show</b>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Enter the following command: <b>ucadmin show</b>

The system displays information about all FC and converged network adapters on the system.

4. Compare the FC settings on the new nodes with the settings that you captured earlier from the original node.
5. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	Modify the FC ports on node4 as needed by entering one of the following commands: <ul style="list-style-type: none"><li>• To program target ports: <b>system node hardware unified-connect modify -type   -t target -adapter port_name</b></li><li>• To program initiator ports: <b>system node unified-connect modify -type   -t initiator -adapter port_name</b></li></ul> <p>-type is the FC4 type, target or initiator.</p>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Modify the FC ports on node4 as needed by entering the following command: <b>ucadmin modify -m fc -t initiator -f adapter_port_name</b> <p>-t is the FC4 type, target or initiator.</p> <p><b>Note:</b> The FC ports need to be programmed as initiators.</p>

6. Take one of the following actions:



<b>If the system that you are upgrading...</b>	<b>Then...</b>
Has storage disks	Verify the new settings by entering the following command and examining the output: <b>system node unified-connect show</b>
A V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Verify the new settings by entering the following command and examining the output: <b>ucadmin show</b>

7. Take one of the following actions:

<b>If the default FC settings on the new nodes are...</b>	<b>Then...</b>
The same as the ones you that captured on the original nodes	Proceed to Step 11.
Different from the ones that you captured on the original nodes	Go to Step 8.

8. Exit Maintenance mode by entering the following command:

**halt**

9. After you enter the command, wait until the system stops at the boot environment prompt.

10. Take one of the following actions:

<b>If the system you are upgrading...</b>	<b>Then...</b>
Is a V-Series system or has FlexArray Virtualization software running Data ONTAP 8.3.0 or later	Access Maintenance mode by entering the following command at the boot environment prompt: <b>boot_ontap maint</b>
Is not a V-Series system and does not have FlexArray Virtualization software	Boot node4 by entering the following command at the boot environment prompt: <b>boot_ontap</b>

11. Take one of the following actions:

<b>If the system that you are upgrading...</b>	<b>Then...</b>
Has storage disks	<ul style="list-style-type: none"> <li>Go to <i>Configuring CNA ports on node4</i> if node4 has a CNA card or CNA onboard ports.</li> <li>Skip the section and go to <i>Mapping ports from node2 to node4</i> if node4 does not have a CNA card or CNA onboard ports.</li> </ul>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<ul style="list-style-type: none"> <li>Go to the section <i>Configuring CNA ports node4</i> if node4 has a CNA card or CNA onboard ports.</li> <li>Skip the section <i>Configuring CNA ports node4</i> if node4 does not have a CNA card or CNA onboard ports, return to the section <i>Installing and booting node4</i>, and resume the section at Step 9.</li> </ul>

## Checking and configuring CNA ports on node4

If node4 has onboard CNA ports or a CNA card, you must check the configuration of the ports and configure them, depending on how you want to use the upgraded system.

### Before you begin

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

### About this task

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

CNA ports might be on an adapter or on the controller with the following configurations:

- CNA cards ordered at the same time as the controller are configured before shipment to have the personality you request.
- CNA cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard CNA ports on new controllers are configured (before shipment) to have the personality you request.

However, you should check the configuration of the CNA ports on node4 and change it, if necessary.

**Important:** If your system has storage disks, you enter the commands in this section at the cluster prompt unless directed to enter Maintenance mode. If you have a V-Series system or a system with FlexArray Virtualization software that is connected to storage arrays, you enter commands in this section in Maintenance mode. You must be in Maintenance mode to configure CNA ports.

### Steps

1. Check how the ports are currently configured by entering one of the following commands on node4:

If the system...	Then...
Has storage disks	Enter the following command:  <b>system node hardware unified-connect show</b>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Enter the following command:  <b>ucadmin show</b>

### Example

The system displays output similar to the following examples:

```
cluster1::> system node hardware unified-connect show
Node  Adapter  Mode   Current Current Pending Pending Admin
-----  -----  -----  -----  -----  -----  -----
f-a    0e       fc     initiator -         -         online
f-a    0f       fc     initiator -         -         online
f-a    0g       cna    target   -         -         online
f-a    0h       cna    target   -         -         online
f-b    0e       fc     initiator -         -         online
f-b    0f       fc     initiator -         -         online
f-b    0g       cna    target   -         -         online
f-b    0h       cna    target   -         -         online
8 entries were displayed.
```

```
*> ucadmin show
Node      Adapter  Current  Current  Pending  Pending  Admin
-----  -
f-a      0e       fc       initiator -        -        online
f-a      0f       fc       initiator -        -        online
f-a      0g       cna      target   -        -        online
f-a      0h       cna      target   -        -        online
f-a      0e       fc       initiator -        -        online
f-a      0f       fc       initiator -        -        online
f-a      0g       cna      target   -        -        online
f-a      0h       cna      target   -        -        online
*>
```

- If the current SFP+ module does not match the desired use, replace it with the correct SFP+ module. Contact your NetApp representative to obtain the correct SFP+ module.
- Examine the output of the `system node hardware unified-connect show` or `ucadmin show` command and determine whether the CNA ports have the personality you want.

- Take one of the following actions:

If the CNA ports...	Then...
Do not have the personality that you want	Go to Step 5.
Have the personality that you want	Skip Step 5 through Step 12 and go Step 13.

- Take one of the following actions:

If the system...	Description
Has storage disks and is running Data ONTAP 8.3	Boot node4 and enter Maintenance mode by entering the following command: <b>boot_ontap maint</b>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Go to Step 6. You should already be in Maintenance mode.

- Take one of the following actions:

If you are configuring...	Then...
Ports on a CNA card	Go to Step 7.
Onboard CNA ports	Skip Step 7 and go to Step 8.

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

```
storage disable adapter adapter_name
```

Adapters in target mode are automatically offline in Maintenance mode.

- If the current configuration does not match the desired use, enter the following command to change the configuration as needed:

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

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

**Note:** You need to use FC initiator for tape drives, FlexArray Virtualization systems, and MetroCluster. You need to use the FC target for SAN clients.

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

```
ucadmin show
```

10. Take one of the following actions:

If the system...	Then...
Has storage disks	<p>a. Enter the following command:</p> <pre>halt</pre> <p>The system stops at the boot environment prompt.</p> <p>b. Enter the following command:</p> <pre>boot_ontap</pre>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays and is running Data ONTAP 8.3	<p>Reboot to Maintenance mode by entering the following command:</p> <pre>boot_ontap maint</pre>

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

If the system...	Then...
Has storage disks	<p>Enter the following command:</p> <pre>system node hardware unified-connect show</pre>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<p>Enter the following command:</p> <pre>ucadmin show</pre>

### Example

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

```
cluster1::> system node hardware unified-connect show
Node  Adapter  Current Mode  Current Type  Pending Mode  Pending Type  Admin Status
-----
f-a   1a        fc      initiator -      -      online
f-a   1b        fc      target  -      initiator online
f-a   2a        fc      target  cna    -      online
f-a   2b        fc      target  cna    -      online
4 entries were displayed.
```

```
*> ucadmin show
Node  Adapter  Current Mode  Current Type  Pending Mode  Pending Type  Admin Status
-----
f-a   1a        fc      initiator -      -      online
f-a   1b        fc      target  -      initiator online
f-a   2a        fc      target  cna    -      online
f-a   2b        fc      target  cna    -      online
4 entries were displayed.
*>
```

12. Place any target ports online by entering one of the following commands, once for each port:

If the system...	Then...
Has storage disks	<code>network fcp adapter modify -node node_name -adapter adapter_name -state up</code>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<code>fcp config adapter_name up</code>

13. Cable the port.

14. Take one of the following actions:

If the system...	Then...
Has storage disks	Go to the section <i>Mapping ports from node2 to node4</i> .
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Return to the section <i>Installing and booting node4</i> and resume the section at Step 9.

### Mapping ports from node2 to node4

You need to make sure that the physical ports on node2 map correctly to the physical ports on node4, which will let node4 communicate with other nodes in the cluster and with the network after the upgrade.

#### Before you begin

Capture information about the ports on the new nodes from the *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com). You will use the information later in this section.

#### About this task

The software configuration of node4 must match the physical connectivity of node4 and IP connectivity must be restored before you continue with the upgrade.

Port settings might vary, depending on the model of the nodes. You must make the original node's port and LIF configuration compatible with what you plan the new node's configuration to be. This is because the new node replays the same configuration when it boots, meaning when you boot node4 that Data ONTAP will try to host LIFs on the same ports that were used on node2.

Therefore, if the physical ports on node2 do not map directly to the physical ports on node4, then software configuration changes will be required to restore cluster, management, and network connectivity after the boot. In addition, if the cluster ports on node2 do not directly map to the cluster ports on node4, node4 may not automatically rejoin quorum when it is rebooted until a software configuration change is made to host the cluster LIFs on the correct physical ports.

#### Steps

- Record all the node2 cabling information for node2, the ports, broadcast domains, and IPspaces, in this table:

LIF	Node2 ports	Node2 IPspaces	Node2 broadcast domains	Node4 ports	Node4 IPspaces	Node4 broadcast domains
Cluster 1						
Cluster 2						
Cluster 3						
Cluster 4						
Cluster 5						

LIF	Node2 ports	Node2 IPspaces	Node2 broadcast domains	Node4 ports	Node4 IPspaces	Node4 broadcast domains
Cluster 6						
Node management						
Cluster management						
Data 1						
Data 2						
Data 3						
Data 4						
SAN						
Intercluster port						

See the “Recording node2 information” section for the steps to obtain this information.

2. Record all the cabling information for node4, the ports, broadcast domains, and IPspaces, in the previous table using the same procedure in the “Recording node2 information” section for the steps to obtain this information.
3. Follow these steps to verify if the setup is a two-node switchless cluster:
  - a. Set the privilege level to advanced:

**Example**

```
cluster::> set -privilege advanced

Warning: These advanced commands are potentially dangerous, use them only when
        directed to do so by NetApp personnel.
Do you wish to continue? (y or n): y

cluster::*>
```

- b. Verify if the setup is a two-node switchless cluster:

**network options switchless-cluster show**

**Example**

```
cluster::*> network options switchless-cluster show

Enable Switchless Cluster: false/true
```

The value of this command must match the physical state of the system.

- c. Return to the administration privilege level:

```
cluster::*> set -privilege admin

cluster::>
```

4. Get node4 into quorum by performing the following steps:
  - a. Boot node4. See “Installing and booting node4” to boot the node if you have not already done so.
  - b. Verify that the new cluster ports are in the Cluster broadcast domain:

```
network port show -node node -port port -fields broadcast-domain
```

#### Example

The following example shows that port e0a is in the Cluster domain on node4:

```
cluster::> network port show -node node4 -port e0a -fields broadcast-domain
node      port broadcast-domain
-----
node4     e1a Cluster
```

- c. If the cluster ports are not in the Cluster broadcast-domain, add them with the following command:

```
broadcast-domain add-ports -ip-space Cluster -broadcast-domain Cluster -ports node:port
```

- d. Add the correct ports to the Cluster broadcast domain:

```
network port modify -node -port -ip-space Cluster -mtu 9000
```

#### Example

This example adds Cluster port e1b on node4:

```
network port modify -node node4 -port e1b -ip-space Cluster -mtu 9000
```

**Note:** For a MetroCluster configuration, you may not be able to change the broadcast domain of a port because it is associated with a port hosting the LIF of a sync-destination SVM and see errors similar to, but not restricted to:

```
command failed: This operation is not permitted on a Vserver that is
configured as the destination of a MetroCluster Vserver relationship.
```

Enter the following command from the corresponding sync-source SVM on the remote site to reallocate the sync-destination LIF to an appropriate port:

```
metrocluster vserver resync -vserver vserver_name
```

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

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

- f. Modify the home port of the cluster LIFs:

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

- g. Remove the old ports from the Cluster broadcast domain:

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

#### Example

This command removes port e0d on node4:

```
network port broadcast-domain remove-ports -ip-space Cluster -broadcast-domain Cluster
-ports node4:e0d
```

- h. Verify that node4 has rejoined quorum:

```
cluster show -node node4 -fields health
```

5. Adjust the broadcast domains hosting your cluster LIFs and node-management/cluster-management LIFs. Ensure that each broadcast domain contains the correct ports. A port cannot be moved between broadcast domains if it is hosting or is home to a LIF so you may need to migrate and modify the LIFs as shown in the following steps:
  - a. Display the home port of a LIF:
 

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

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

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

```
network interface modify -vserver vservers -lif lif_name -home-port port_name
```
6. Adjust the intercluster broadcast domains and migrate the intercluster LIFs, if necessary, using the same commands shown in Step 6.
7. Adjust any other broadcast domains and migrate the data LIFs, if necessary, using the same commands shown in Step 5.
8. If there were any ports on node2 that no longer exist on node4, follow these steps to delete them:

- a. Access the advanced privilege level on either node:

```
set -privilege advanced
```

The system displays the following message:

```
Warning: These advanced commands are potentially dangerous; use them only when
directed to do so by NetApp personnel.
do you wish to continue? (y or n):
```

- b. Enter **y**.
- c. To delete the ports:

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

- d. Return to the admin level:

```
set -privilege admin
```

9. Adjust all the LIF failover groups:

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

### Example

The following command sets the failover policy to `broadcast-domain-wide` and uses the ports in failover group `fg1` as failover targets for LIF `data1` on node4:

```
network interface modify -vserver node4 -lif data1 failover-policy broadcast-domain-wide -
failover-group fg1
```

See “Configuring failover settings on a LIF” in the *ONTAP 9 Network Management Guide* or *ONTAP 9 Commands: Manual Page Reference* for more information.

10. Verify the changes on node4:

```
network port show -node node4
```



## Verifying the node4 installation

After you install and boot node4, you need to verify that it is installed correctly, that it is part of the cluster, and that it can communicate with node3.

### Steps

1. At the system prompt, log in to node4.
2. Verify that node4 is both part of the same cluster as node3 and healthy by entering the following command:  
`cluster show`
3. Verify that node4 can communicate with node3 and that all LIFs are up by entering the following command:  
`network interface show -curr-node node4`

4. Take one of the following actions:

If node4 is...	Then...
In a chassis separate from node3	Connect the HA interconnect between the nodes by completing the following steps: <ol style="list-style-type: none"><li>a. Connect the top interconnect port of node3 to the top interconnect port of node4.</li><li>b. Connect the bottom interconnect port of node3 to the bottom interconnect port of node4.</li><li>c. Go to Step 5.</li></ol>
In the same chassis as node3	Go to Step 5.  You do not need to manually connect the HA interconnect between the nodes; in same-chassis configurations, the HA interconnect is connected automatically through the backplane.

5. Take one of the following actions:

If the cluster is...	Then...
In a SAN environment	Complete Step 6 and go to the section <i>Moving non-SAN data LIFs owned by node2 from node3 to node4 and verifying SAN LIFs on node4.</i>
Not in a SAN environment	Skip Step 6 go to the section <i>Moving non-SAN data LIFs owned by node2 from node3 to node4 and verifying SAN LIFs on node4.</i>

6. Verify that both node3 and node4 are in quorum by entering the following command on one of the nodes:

```
event log show -messagename scsiblade.*
```

### Example

The following example shows the output when the nodes in the cluster are in quorum:

```
cluster::> event log show -messagename scsiblade.*
Time           Node           Severity      Event
-----
8/13/2012 14:03:51 node1          INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2012 14:03:51 node2          INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2012 14:03:48 node3          INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2012 14:03:43 node4          INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
```

## Moving non-SAN data LIFs owned by node2 from node3 to node4 and verifying SAN LIFs on node4

After you verify the node4 installation and before you relocate node2 aggregates from node3 to node4, you need to move the non-SAN data LIFs owned by node2 currently on node3 from node3 to node4. You also need to verify the SAN LIFs on node4.

### About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. Moving SAN LIFs is not necessary for cluster or service health during the upgrade. SAN LIFs are not moved unless they need to be mapped to new ports. You verify that the LIFs are healthy and located on appropriate ports after you bring node4 online.

### Steps

1. List all the non-SAN data LIFs that are not owned by node3 by entering the following command on either node and capturing the output:

```
network interface show -role data -curr-node node3 -is-home false
```

2. Take one of the following actions:

If node2...	Description
Had interface groups or VLANs configured	Go to Step 3.
Did not have interface groups or VLANs configured	Skip Step 3 and go to Step 4.

3. Take the following steps to migrate any non-SAN data LIFs hosted on interface groups and VLANs that originally were on node2 from node3 to node4.

- a. Migrate any LIFs hosted on node3 that previously belonging to node2 on an interface group to a port on node4 that is capable of hosting LIFs on the same network by entering the following command, once for each LIF:

```
network interface migrate -vserver vs_server_name -lif lif_name -destination-node node4 -destination-port netport|ifgrp
```

- b. Modify the home port and home node of the LIFs in substep a to the port and node currently hosting the LIFs by entering the following command, once for each LIF:

```
network interface modify -vserver vs_server_name -lif data_lif_name -home-node node4 -home-port netport|ifgrp
```

- c. Migrate any LIFs hosted on node3 that previously belonged to node2 on a VLAN port to a port on node4 that is capable of hosting LIFs on the same network by entering the following command, once for each LIF:

```
network interface migrate -vserver vs_server_name -lif data_lif_name -destination-node node4 -destination-port netport|ifgrp
```

- d. Modify the home port and home node of the LIFs in substep c to the port and node currently hosting the LIFs by entering the following command, once for each LIF:

```
network interface modify -vserver vs_server_name -lif data_lif_name -home-node node4 -home-port netport|ifgrp
```

4. Take one of the following actions:

If the cluster is configured for...	Then...
NAS	Complete Step 5 through Step 8, skip Step 9, and complete Step 10 through Step 13.
SAN	Skip Step 5 through Step 8, and complete Step 9 through Step 13.
Both NAS and SAN	Complete Step 5 through Step 13.

- If you have data ports that are not the same on your platforms, enter the following command to add the ports to the broadcast domain:

```
network port broadcast-domain add-ports -ipSPACE IPspace_name -broadcast-domain mgmt -ports node:port
```

#### Example

The following example adds port e0a on node 6280-1 and port e0i on node 8060-1 to broadcast domain mgmt in the IPspace Default:

```
cluster::> network port broadcast-domain add-ports -ipSPACE Default
-broadcast-domain mgmt -ports 6280-1:e0a, 8060-1:e0i
```

- Migrate each non-SAN data LIF to node4 by entering the following command, once for each LIF:

```
network interface migrate -vserver vservice_name -lif datalif_name -destination-node node4 -
destination-port netport|ifgrp -home-node node4
```

- Make sure that the data migration is persistent by entering the following command:

```
network interface modify -vserver vservice_name -lif datalif_name -home-port netport|ifgrp
```

- Verify the status of all links as up by entering the following command to list all the network ports and examining its output:

```
network port show
```

#### Example

The following example shows the output of the `network port show` command with some LIFs up and others down:

```
cluster::> network port show
```

Node	Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Oper
node3						
	a0a	Default	-	up	1500	auto/1000
	e0M	Default	172.17.178.19/24	up	1500	auto/100
	e0a	Default	-	up	1500	auto/1000
	e0a-1	Default	172.17.178.19/24	up	1500	auto/1000
	e0b	Default	-	up	1500	auto/1000
	e1a	Cluster	Cluster	up	9000	auto/10000
	e1b	Cluster	Cluster	up	9000	auto/10000
node4						
	e0M	Default	172.17.178.19/24	up	1500	auto/100
	e0a	Default	172.17.178.19/24	up	1500	auto/1000
	e0b	Default	-	up	1500	auto/1000
	e1a	Cluster	Cluster	up	9000	auto/10000
	e1b	Cluster	Cluster	up	9000	auto/10000

12 entries were displayed.

- If the output of the `network port show` command displays network ports that are not available in the new node and are present in the old nodes, delete the old network ports by completing the following substeps:

- Enter the advanced privilege level by entering the following command:

```
set -privilege advanced
```

The system displays the following message:

```
Warning: These advanced commands are potentially dangerous; use them only when directed
to do so by NetApp personnel.
do you wish to continue? (y or n):
```

- b. Enter `y`.
- c. Enter the following command, once for each old network port:  
`network port delete -node node_name -port port_name`
- d. Return to the admin level by entering the following command:  
`set -privilege admin`

**10.** Ensure that the SAN LIFs are on the correct ports on node4 by completing the following substeps:

- a. Enter the following command:  
`network interface show -data-protocol iscsi|fc -home-node node4`

**Example**

The system returns output similar to the following example:

```
cluster::> network interface show -data-protocol iscsi|fc -home-node node4
-----
Vserver      Logical      Status      Network      Current      Current      Is
Interface    Admin/Oper   Address/Mask Node          Port         Home
-----
vs0
  a0a         up/down    10.63.0.53/24  node3        a0a         true
  data1       up/up      10.63.0.50/18  node3        e0c         true
  rads1       up/up      10.63.0.51/18  node3        e1a         true
  rads2       up/down    10.63.0.52/24  node3        e1b         true
vs1
  lif1        up/up      172.17.176.120/24 node3        e0c         true
  lif2        up/up      172.17.176.121/24 node3        e1a         true
```

- b. If node4 has any SAN LIFs or groups of SAN LIFs that are on a port that did not exist on node2, move them to an appropriate port on node4 by entering one of the following commands:

- i. Set the LIF status to down by entering the following command:  
`network interface modify -vserver vsserver_name -lif lif_name -status-admin down`

- ii. Remove the LIF from the port set:  
`portset remove -vserver vsserver_name -portset portset_name -port-name port_name`

- iii. Enter one of the following commands:

- To move a single LIF, enter the following command:  
`network interface modify -lif lif_name -home-port new_home_port`
- To move all the LIFs on a single nonexistent or incorrect port to a new port, enter the following command:  
`network interface modify {-home-port port_on_node2 -home-node node2 -role data} -home-port new_home_port_on_node4`
- Add the LIFs back to the port set:  
`portset add -vserver vsserver_name -portset portset_name -port-name port_name`

**Note:** You need to ensure that you move SAN LIFs to a port that has the same link speed as the original port.

**11.** Modify the status of all LIFs to **up** so the LIFs can accept and send traffic on the node by entering the following command:

```
network interface modify -vserver vsserver_name -home-port port_name -home-node node4 -lif lif_name -status-admin up
```

**12.** Verify that any SAN LIFs have been moved to the correct ports and that the LIFs have the status of up by entering the following command on either node and examining the output:

```
network interface show -home-node node4 -role data
```

13. 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
```

## Relocating node2's non-root aggregates from node3 to node4

Having relocated node2's non-root aggregates to node3, you now must relocate them from node3 to node4.

### Steps

1. Enter the following command on either controller and examine the output to identify which non-root aggregates to relocate:

```
storage aggregate show -owner-name node3 -home-id node2_system_id
```

2. Relocate the aggregates by completing the following substeps:

- a. Access the advanced privilege level by entering the following command on either node:

```
set -privilege advanced
```

The system displays the following message:

```
Warning: These advanced commands are potentially dangerous; use them only when
directed to do so by NetApp personnel.
do you wish to continue? {y or n}:
```

- b. Enter **y**.

- c. Enter the following command:

```
storage aggregate relocation start -node node3 -destination node4 -aggregate-list
aggr_name1, aggr_name2... -ndo-controller-upgrade true
```

The aggregate list is the list of aggregates owned by node4 that you obtained in Step 1.

The system displays the following message:

```
Warning: Aggregate relocation will not change home ownership of an aggregate
which is owned by the source node during an NDO controller upgrade. This
parameter should be used only while performing NDO controller upgrade.
Are you performing NDO controller upgrade? {y|n}:
```

- d. When prompted, enter **y**.

Relocation occurs in the background. It could take anywhere from a few seconds to a couple of minutes to relocate an aggregate. The time includes both client outage and non-outage portions. The command does not relocate any offline or restricted aggregates.

- e. Return to the admin level by entering the following command:

```
set -privilege admin
```

3. Check the relocation status by entering the following command:

```
storage aggregate relocation show -node node3
```

The output will display Done for an aggregate after it has been relocated.

**Note:** Wait until all the node2 aggregates have been relocated to node4 before proceeding to the next step.

4. Take one of the following actions:

If relocation of...	Then...
All aggregates was successful	Go to Step 5.
Any aggregates failed, or were vetoed	<p data-bbox="609 279 1511 306">a. Check the EMS logs for the corrective action.</p> <p data-bbox="609 327 1511 354">b. Perform the corrective action.</p> <p data-bbox="609 375 1511 485">c. Access the advanced privilege level by entering the following command on either node:  <b>set -privilege advanced</b>  The system displays the following message:</p> <div data-bbox="649 512 1520 642" style="background-color: #f0f0f0; padding: 5px;"> <pre data-bbox="649 512 1520 642">Warning: These advanced commands are potentially dangerous; use them only when directed to do so by NetApp personnel. do you wish to continue? {y or n}:</pre> </div> <p data-bbox="609 663 1511 690">d. Enter <b>y</b>.</p> <p data-bbox="609 711 1511 911">e. Relocate any failed or vetoed aggregates by entering the following command:  <b>storage aggregate relocation start -node node3 -destination node4 -aggregate-list aggr_name1, aggr_name2... ndo-controller-upgrade true</b>  The aggregate list is the list of failed or vetoed aggregates.  The system displays the following message:</p> <div data-bbox="649 938 1520 1142" style="background-color: #f0f0f0; padding: 5px;"> <pre data-bbox="649 938 1520 1142">Warning: Aggregate relocation will not change home ownership of an aggregate which is owned by the source node during an NDO controller upgrade. This parameter should be used only while performing NDO controller upgrade. Are you performing NDO controller upgrade? {y n}:</pre> </div> <p data-bbox="609 1163 1511 1190">f. When prompted, enter <b>y</b>.</p> <p data-bbox="609 1211 1511 1289">g. Return to the admin level by entering the following command:  <b>set -privilege admin</b></p> <p data-bbox="609 1310 1511 1337">If necessary, you can force the relocation using one of the following methods:</p> <ul data-bbox="609 1358 1511 1528" style="list-style-type: none"> <li data-bbox="609 1358 1511 1415">• Overriding veto checks by using the command <code>storage aggregate relocation start -override-vetoes -ndo-controller-upgrade</code></li> <li data-bbox="609 1436 1511 1528">• Overriding destination checks by using the command <code>storage aggregate relocation start -override-destination-checks -ndo-controller-upgrade</code></li> </ul> <p data-bbox="609 1549 1511 1608">See the <i>ONTAP 9 Disks and Aggregates Power Guide</i> and the <i>ONTAP 9 Commands: Manual Page Reference</i> for more information about storage aggregate relocation commands.</p>

5. Verify that all node2 non-root aggregates are online and their state on node4 by entering the following command:

```
storage aggregate show -node node4 -state offline -root false
```

The node2 aggregates were listed in the output of the command in Step 1.

6. If any aggregate has gone offline or become foreign, bring it online by using the following command for each aggregate:

```
storage aggregate online -aggregate aggr_name
```

7. Verify that all the volumes in node2 aggregates are online on node4 by entering the following command:

```
volume show -node node4 -state offline
```

8. If any volumes are offline on node4, bring them online by entering the following command:

```
volume online -vserver vservice-name -volume volume_name
```

9. Send a post-upgrade AutoSupport message to NetApp for node4 by entering the following command:

```
system node autosupport invoke -node node4 -type all -message "node2 successfully upgraded from platform_old to platform_new"
```

## Upgrading the node pair: Stage 6

During stage 6, ensure that the new nodes are set up correctly. If one of the new nodes has a unified target adapter, you must restore any port configurations and might need to change the personality of the adapter. You also should set up Storage Encryption if the new nodes are encryption-enabled. You also should decommission the old nodes.

### Steps

1. [Ensuring that the new controllers are set up correctly](#) on page 87
2. [Setting up Storage Encryption on the new controller module](#) on page 90
3. [Setting up Storage Encryption on the new nodes running Data ONTAP 8.3 or earlier](#) on page 90
4. [Decommissioning the old system](#) on page 91
5. [Resuming SnapMirror operations](#) on page 92

## Ensuring that the new controllers are set up correctly

To ensure correct setup, you must enable the HA pair. You must also verify that node3 and node4 can access each other's storage and that neither owns data LIFs belonging to other nodes on the cluster. In addition, you must ensure that node3 owns node1's aggregates and that node4 owns node2's aggregates, and that the volumes for both nodes are online.

### Steps

1. Enable storage failover by entering the following command on one of the nodes:

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

2. Verify that storage failover is enabled by entering the following command:

```
storage failover show
```

### Example

The following example shows the output of the command when storage failover is enabled:

```
cluster::> storage failover show
Node           Partner           Takeover
-----           -
node3          node4             true    Connected to node4
node4          node3             true    Connected to node3
```

3. Take one of the following actions:

If the cluster is a...	Description
Two-node cluster	Enable cluster high availability by entering the following command on either node: <pre>cluster ha modify -configured true</pre>
Cluster with more than two nodes	Go to Step 4.

4. Verify that node3 and node4 belong to the same cluster by entering the following command and examining the output:

```
cluster show
```

5. Verify that node3 and node4 can access each other's storage by entering the following command and examining the output:

```
storage failover show -fields local-missing-disks,partner-missing-disks
```

6. Verify that neither node3 nor node4 owns data LIFs home-owned by other nodes in the cluster by entering the following command and examining the output:

```
network interface show
```

If either node3 or node4 owns data LIFs home-owned by other nodes in the cluster, use the `network interface revert` command to revert the data LIFs to their home-owner.

7. Verify that node3 owns the aggregates from node1 and that node4 owns the aggregates from node2 by entering the following commands and comparing that output with the output you collected in substep 18c of the section *Preparing the nodes for the upgrade*:

```
storage aggregate show -owner-name node3
```

```
storage aggregate show -owner-name node4
```

8. Determine whether any volumes are offline by entering the following commands:

```
volume show -node node3 -state offline
```

```
volume show -node node4 -state offline
```

9. If any volumes are offline, compare them with the list of offline volumes that you captured in substep 18d of the section *Preparing the nodes for the upgrade*, and bring any online any of the offline volumes, as required, by entering the following command, once for each volume:

```
volume online -vserver vsserver-name -volume volume_name
```

10. Install new licenses for the new nodes by entering the following commands for each node:

```
system license add -license-code license_code,license_code,license_code...
```

The `license-code` parameter accepts a list of 28 upper-case alphabetic character keys. You can add one license at a time, or you can add or multiple licenses at once, each license key separated by a comma.

11. To remove all of the old licenses from the original nodes, enter one of the following commands:

```
system license clean-up -unused -expired
```

```
system license delete -serial-number node_serial_number -package licensable_package
```

- To delete all expired licenses, enter:

```
system license clean-up -expired
```

- To delete all unused licenses, enter:

```
system license clean-up -unused
```

- To delete a specific license from a cluster, enter the following commands on the nodes:

```
system license delete -serial-number <node1 serial number> -package *
```

```
system license delete -serial-number <node2 serial number> -package *
```

The following output is displayed:

```
Warning: The following licenses will be removed:
<list of each installed package>
Do you want to continue? {y|n}: y
```

Enter `y` to remove all of the packages.

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



**system license show**

You might want to compare the output with the output that you captured in Step 29 of the “Preparing the nodes for the upgrade” section.

13. Configure the SPs by performing the following command on both nodes:

**system service-processor network modify -node node\_name**

See the *System Administration Reference* for information about the SPs and the *ONTAP 9 Commands: Manual Page Reference* for detailed information about the `system node service-processor network modify` command.

14. Take one of the following actions:

If the upgraded system is...	Then...
A 62xx or 80xx system and the system is in a CIFS environment	Go to Step 15.
Any system except a 62xx or 80xx system, regardless of the environment	Skip Step 15 and Step 16 and go to Step 17.

15. Take one of the following actions:

If the 62xx or 80xx system...	Then...
Will have workloads in which less than 50% of the operations will be 4 KB or smaller	Skip Step 16 and go to Step 17.
Will have workloads in which more than 50% of the operations will be 4 KB or smaller	Go to Step 16.

16. Take the following actions on one of the new nodes:

- a. Enter advanced privilege level by entering the following command:

**set -privilege advanced**

The system displays the following message:

```
Warning: These advanced commands are potentially dangerous; use them only when
directed to do so by NetApp personnel.
do you wish to continue? (y or n):
```

- b. Enter **y**.

- c. Enter the following command:

**storage failover modify -node node\_name -cifs-ndo-duration default|medium|low**

- Enter **medium** if the system will have workloads in which 50 percent to 75 percent of the operations will be 4 KB or smaller.
- Enter **low** if the system will have workloads in which 75 percent to 100 percent of the operations will be 4 KB or smaller.

- d. Return to the admin level by entering the following command:

**set -privilege admin**

- e. Reboot the system to ensure that the changes take effect.

17. If you want to set up a switchless cluster on the new nodes, follow the instructions in *Transitioning to a two-node switchless cluster* on the NetApp Support Site.

## After you finish

If Storage Encryption is enabled on node3 and node4, complete the section *Setting up Storage Encryption on the new nodes*. Otherwise, complete the section *Decommissioning the old system*.

## Setting up Storage Encryption on the new controller module

If the replaced controller or the HA partner of the new controller uses Storage Encryption, you must configure the new controller module for Storage Encryption, including installing SSL certificates and setting up key management servers.

### About this task

This procedure includes steps that are performed on the new controller module. Be sure to enter the command on the correct node.

### Steps

1. Verify that the key management servers are still available, their status, and their authentication key information:

```
security key-manager show -status
```

```
security key-manager query
```

2. Add the key management servers listed in step 1 to the key management server list in the new controller:

- a. Add the key management server:

```
security key-manager -add key_management_server_ip_address
```

- b. Repeat the previous step for each listed key management server.

You can link up to four key management servers.

- c. Verify that the key management servers were added successfully:

```
security key-manager show
```

3. On the new controller module, run the key management setup wizard to set up and install the key management servers.

You must install the same key management servers that are installed on the existing controller module.

- a. Launch the key management server setup wizard on the new node:

```
security key-manager setup -node new_controller_name
```

- b. Complete the steps in the wizard to configure key management servers.

4. Restore authentication keys from all linked key management servers to the new node:

```
security key-manager restore -node new_controller_name
```

## Setting up Storage Encryption on the new nodes running Data ONTAP 8.3 or earlier

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

### Before you begin

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

### About this task

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

If you used Storage Encryption on the original system and migrated the disk shelves to the new system, you can reuse the SSL certificates that are stored on migrated disk drives for Storage Encryption functionality on the upgraded system. However, you

should check that the SSL certificates are present on the migrated disk drives. If they are not present you will need to obtain them.

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

### Steps

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

Requirements for obtaining the certificates and instructions for installing them are contained in the *Clustered Data ONTAP Software Setup Guide*.

2. Collect the information required to configure Storage Encryption on the new nodes.

This includes the network interface name, the network interface IP address, and the IP address for external key management server. The required information is contained in the *Clustered Data ONTAP Software Setup Guide*.

3. Launch and run the Storage Encryption setup wizard, responding to the prompts as appropriate.
4. If you have not done so, repeat Step 1 through Step 3 on the other new node.

### After you finish

See the *ONTAP 9 Disks and Aggregates Power Guide* for information about managing Storage Encryption on the updated system.

## Decommissioning the old system

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

### Steps

1. Go to the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com) and log in.
2. Click the link **My Installed Systems**.
3. On the **Installed Systems** page, enter the serial number of the old system in the form and then click **Go!**  
A new page displays information about the controller.
4. Make sure that the information about the controller is correct.

---

<b>If the information about the controller is...</b>	<b>Then...</b>
--	----------------

---

Correct...

- a. Select **Decommission this system** in the Product Tool Site drop-down menu.
- b. Go to Step 5.

---

Not correct...

- a. Click the feedback link to open the form for reporting the problem.
  - b. Fill out and submit the form.
- 

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

## Resuming SnapMirror operations

You can upgrade the cross cluster asynchronous disaster recovery (ADR) and resume with the Storage Virtual Machine (SVM, formerly known as Vserver) disaster recovery relationship. The updates are on schedule once the upgrade is completed.

### Steps

1. Verify the Storage Virtual Machine (SVM) disaster recovery SnapMirror state on the destination.

#### Example

```
cluster::*> snapmirror show
```

2. Resume the SnapMirror relationship:

#### Example

```
cluster::*> snapmirror resume -destination-vserver
```

## Troubleshooting

You might encounter a failure while upgrading the node pair. The node might crash, aggregates might not relocate, or LIFs might not migrate. The cause of the failure and its solution depend on when the failure occurred during the upgrade procedure.

Refer to the illustration and the table describing the different phases of the procedure in the section *Overview of the upgrade using ARL*. The information about failures that can occur is listed by the phase of the procedure.

## ARL failures

ARL might fail at different points during the upgrade.

### Checking for aggregate relocation failure

During the procedure, ARL might fail in Stage 2, Stage 3, or Stage 5.

#### Steps

1. Enter the following command: `storage aggregate relocation show` and examine its output.

The `storage aggregate relocation show` command shows you which aggregates were successfully relocated and which ones were not, along with the causes of failure.

2. Check the console for any EMS messages.

3. Take one of the following actions:

- Take the appropriate corrective action, depending on the output of the `storage aggregate relocation show` command and the output of the EMS message.
- Force relocation of the aggregate or aggregates by using the `override-vetoes` option or the `override-destination-checks` option of the `storage aggregate relocation start` command.

For detailed information about the `storage aggregate relocation start`, `override-vetoes`, and `override-destination-checks` options, see the *ONTAP 9 Commands: Manual Page Reference*.

## Aggregates originally on node1 are owned by node 4 after completion of the upgrade

At the end of the upgrade procedure, node3 should be the new home node of aggregates that originally had node1 as the home node. You can relocate them after the upgrade.

### About this task

Aggregates might fail to relocate properly, having node1 as their home node instead of node3 under the following circumstances:

- During Stage 3, when aggregates are relocated from node2 to node3. Some of the aggregates being relocated have node1 as their home node. For example, such an aggregate could be called `aggr_node_1`. If relocation of `aggr_node_1` fails during Stage 3, and relocation cannot be forced, then the aggregate will be left behind on node2.
- After Stage 4, when node2 is replaced with node4. When node2 is replaced, `aggr_node_1` will come online with node4 as its home node instead of node3.

You can fix the incorrect ownership problem after Stage 6 once storage failover has been enabled by completing the following steps:

### Steps

1. Enter the following command to get a list of aggregates:

```
storage aggregate show -nodes node4 -is-home true
```

To identify aggregates that were not correctly relocated, refer to the list of aggregates with the home owner of node1 that you obtained in the section *Preparing the nodes for the upgrade* and compare it with output of the above command.

2. Compare the output of Step 1 with the output you captured for node1 in the section *Preparing the nodes for the upgrade* and note any aggregates that were not correctly relocated.

3. Relocate the aggregates left behind on node4 by entering the following command:

```
storage aggregate relocation start -node node4 -aggr aggr_node_1 -destination node3
```

Do not use the `-ndo-controller-upgrade` parameter during this relocation.

4. Enter the following command to verify that node3 is now the home owner of the aggregates:

```
storage aggregate show -aggregate aggr1,aggr2,aggr3... -fields home-name
```

`aggr1,aggr2,aggr3...` is the list of aggregates that had node1 as the original home owner.

Aggregates that do not have node3 as home owner can be relocated to node3 using the same relocation command in Step 3.

## Reboots, panics, or power cycles

The system might crash – reboot, panic or go through a power cycle – during different stages of the upgrade. The solution to these problems depends on when they occur.

### Reboots, panics, or power cycles during Stage 2

Crashes can occur before, during, or immediately after Stage 2, during which you relocate aggregates from node1 to node2, move data LIFs and SAN LIFs owned by node1 to node2, record node1 information, and retire node1.

## Node1 or node2 crashes before Stage 2 with HA still enabled

If either node1 or node2 crashes before Stage 2, no aggregates have been relocated yet and the HA configuration is still enabled.

### About this task

Takeover and giveback can proceed normally.

### Steps

1. Check the console for EMS messages that the system might have issued, and take the recommended corrective action.
2. Continue with the node-pair upgrade procedure.

## Node1 crashes during or just after Stage 2 with HA still enabled

Some or all aggregates have been relocated from node1 to node2, and HA is still enabled. Node2 will take over node1's root volume and any non-root aggregates that were not relocated.

### About this task

Ownership of aggregates that were relocated will look the same as the ownership of non-root aggregates that were taken over because home owner has not changed.

When node1 enters the `waiting for giveback` state, node2 will give back all the node1 non-root aggregates.

### Steps

1. Complete Step 1 in the section *Relocating non-root aggregates from node1 to node2* again.
2. Continue with the node-pair upgrade procedure.

## Node1 crashes after Stage2 while HA is disabled

Node2 will not take over but it is still serving data from all non-root aggregates.

### Steps

1. Bring up node1.
2. Continue with the node-pair upgrade procedure.

You might see some changes in the output of the `storage failover show` command, but that is typical and does not affect the procedure. See the troubleshooting section *Unexpected storage failover show command output*.

## Node2 crashes during or after Stage 2 with HA still enabled

Node1 has relocated some or all of its aggregates to node2. HA is enabled.

### About this task

Node1 will take over all of node2's aggregates as well any of its own aggregates that it had relocated to node2. When node2 enters the `waiting for Giveback` state, node1 will give back all of node2's aggregates.

### Steps

1. Complete Step 1 in the section *Relocating non-root aggregates from node1 to node2* again.
2. Continue with the node-pair upgrade procedure.

### **Node2 crashes after Stage 2 and after HA is disabled**

Node1 will not take over.

#### **Steps**

1. Bring up node2.  
A client outage will occur for all aggregates while node2 is booting up.
2. Continue with rest of the node pair upgrade procedure.

### **Reboots, panics, or power cycles during Stage 3**

Crashes can occur during or immediately after Stage 3, during which you install and boot node3, map ports from node1 to node3, move data LIFs and SAN LIFs belonging to node1 and node2 to node3, and relocate all aggregates from node2 to node3.

### **Node2 crash during Stage 3 with HA disabled and before relocating any aggregates**

Node3 will not take over following a node2 crash as HA is already disabled.

#### **Steps**

1. Bring up node2.  
A client outage will occur for all aggregates while node2 is booting up.
2. Continue with the node-pair upgrade procedure.

### **Node2 crashes during Stage 3 after relocating some or all aggregates**

Node2 has relocated some or all of its aggregates to node3, which will serve data from aggregates that were relocated. HA is disabled.

#### **About this task**

There will be client outage for aggregates that were not relocated.

#### **Steps**

1. Bring up node2.
2. Relocate the remaining aggregates by completing Step 1 through Step 3 in the section *Relocating aggregates from node2 to node3*.
3. Continue with the node-pair upgrade procedure.

### **Node3 crashes during Stage 3 and before node2 has relocated any aggregates**

Node2 will not take over but it is still serving data from all non-root aggregates.

#### **Steps**

1. Bring up node3.
2. Continue with the node-pair upgrade procedure.

### **Node3 crashes during Stage 3 during aggregate relocation**

If node3 crashes while node2 is relocating aggregates to node3, node2 will abort the relocation of any remaining aggregates.

#### **About this task**

Node2 continues to serve remaining aggregates, but aggregates that were already relocated to node3 encounter client outage while node3 is booting.

#### **Steps**

1. Bring up node3.
2. Complete Step 3 again in the section *Relocating all aggregates from node2 to node3*.
3. Continue with the node-pair upgrade procedure.

### **Node3 fails to boot after crashing in Stage 3**

Because of a catastrophic failure, node3 cannot be booted following a crash during Stage 3.

#### **Step**

1. Contact technical support.

### **Node2 crashes after Stage 3 but before Stage 5**

Node3 continues to serve data for all aggregates. The HA pair is disabled.

#### **Steps**

1. Bring up node2.
2. Continue with the node-pair upgrade procedure.

### **Node3 crashes after Stage 3 but before Stage 5**

Node3 crashes after Stage 3 but before Stage 5. The HA pair is disabled.

#### **Steps**

1. Bring up node3.  
There will be a client outage for all aggregates.
2. Continue with the node-pair upgrade procedure.

### **Reboots, panics, or power cycles during Stage 5**

Crashes can occur during Stage 5, the stage in which you install and boot node4, map ports from node2 to node4, move data LIFs and SAN LIFs belonging to node2 from node3 to node4, and relocate all of node2's aggregates from node3 to node4.

### **Node3 crashes during Stage 5**

Node3 has relocated some or all of node2's aggregates to node4. Node4 does not take over but continues to serve non-root aggregates that node3 already relocated. The HA pair is disabled.

#### **About this task**

There is be an outage for the rest of the aggregates until node3 boots again.



## Steps

1. Bring up node3.
2. Relocate the remaining aggregates that belonged to node2 by repeating Step 1 through Step 3 in the section *Relocating node2's non-root aggregates from node3 to node4*.
3. Continue with the node pair upgrade procedure.

## Node4 crashes during Stage 5

Node3 has relocated some or all of node2's aggregates to node4. Node3 does not take over but continues to serve non-root aggregates that node3 owns as well as those that were not relocated. HA is disabled.

### About this task

There is an outage for non-root aggregates that were already relocated until node4 boots again.

## Steps

1. Bring up node4.
2. Relocate the remaining aggregates that belonged to node2 by again completing Step 1 through Step 3 in *Relocating node2's non-root aggregates from node3 to node4*.
3. Continue with the node-pair upgrade procedure.

## Issues that can arise in multiple stages of the procedure

Some issues can occur in different stages of the procedure.

### Unexpected storage failover show command output

During the procedure, if the node that hosts all data aggregates panics or is rebooted accidentally, you might see unexpected output for the `storage failover show` command before and after the reboot, panic, or power cycle.

### About this task

You might see unexpected output from the `storage failover show` command in Stage 2, Stage 3, Stage 4, or Stage 5.

The following example shows the expected output of the `storage failover show` command if there are no reboots or panics on the node that hosts all the data aggregates:

```
cluster::> storage failover show
Node           Partner           Takeover
-----
node1          node2             false
node2          node1             false
State Description
-----
Unknown
Node owns partner aggregates as part
of the non-disruptive head upgrade
procedure., Takeover is not
possible: Storage failover is
disabled

2 entries were displayed.
```

The following example shows the output of the `storage failover show` command after a reboot or panic:

```
cluster::> storage failover show
Node           Partner           Takeover
-----
node1          node2             -
State Description
-----
Unknown
```

```
node2          node1          false      Waiting for node1, Partial
              giveback, Takeover is not possible:
              Storage failover is disabled

2 entries were displayed.
```

Although the output says that a node is in partial giveback and that storage failover is disabled, you can disregard this message.

### Step

1. No action is required; continue with the node-pair upgrade procedure.

## LIF migration failure

After you migrate them, LIFs might not come online after migration in Stage 2, Stage 3, or Stage 5.

### About this task

### Steps

1. Enter the following command:

```
network port show -node node_name -type physical
```

`node_name` is the node to which you migrated LIFs.

### Example

The following example shows the output of the command for a node named `node1`:

```
cluster::> network port show -node node1 -type physical

Node  Port  IPspace      Broadcast Domain Link  MTU      Speed (Mbps)
-----  ---  -
node1
      e0M   Default      Default      up    1500     auto/1000
      e0a   Default      Default      down  1500     auto/10
      e0b   Default      Default      down  1500     auto/10000
      e0c   Cluster     Cluster     up    9000     auto/10000
      e0d   Cluster     Cluster     up    9000     auto/10000
      e0i   Default      Default      up    1500     auto/1000
      e0j   Default      Default      up    1500     auto/1000
      e0k   Default      Default      down  1500     auto/10
      e0l   Default      Default      down  1500     auto/10

9 entries were displayed.
```

2. Examine the output of the command in Step 1 and verify that the IPspace results match the source node. Use the output that you captured for the original node in the section *Recording node1 information* or *recording node2 information*.
3. Ensure that port MTU size is the same as that of the source node. For example, if the cluster port MTU size is 9000 on the source node, it should be 9000 on the destination node.
4. Check the physical connectivity of the network cable if the physical state of the port is down.

# How to send comments about documentation and receive update notifications

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

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

[doccomments@netapp.com](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

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

You can also contact us in the following ways:

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

## Trademark information

Active IQ, AltaVault, Arch Design, ASUP, AutoSupport, Campaign Express, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Element, Fitness, Flash Accel, Flash Cache, Flash Pool, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, Fueled by SolidFire, GetSuccessful, Helix Design, LockVault, Manage ONTAP, MetroCluster, MultiStore, NetApp, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANscreen, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, SolidFire Helix, StorageGRID, SyncMirror, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web.

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