



Upgrade by using ARL

Upgrade controllers

NetApp

February 10, 2026

Table of Contents

Upgrade by using ARL	1
Start here: Choose your ARL upgrade procedure	1
Use "system controller replace" commands with ONTAP 9.15.1 or later	1
Upgrade controller models in the same chassis	1
Use "system controller replace" commands with ONTAP 9.8 or later	2
Use "system controller replace" commands with ONTAP 9.5 to 9.7	3
Use manual ARL commands	3
Use "system controller replace" commands to upgrade controller hardware introduced in ONTAP 9.15.1 and later	4
Learn about this ARL upgrade procedure	4
Automate the controller upgrade process	6
Decide whether to use this aggregate relocation procedure	6
Required tools and documentation	8
Guidelines for upgrading controllers with ARL	8
Learn about the ARL upgrade sequence	9
Stage 1. Prepare for upgrade	11
Stage 2. Relocate and retire node1	16
Stage 3. Install and boot node3	19
Stage 4. Relocate and retire node2	39
Stage 5. Install and boot node4	41
Stage 6. Complete the upgrade	61
Troubleshoot	67
References	74
Use "system controller replace" commands to upgrade controller models in the same chassis	75
Learn about this ARL upgrade procedure	75
Decide whether to use this aggregate relocation procedure	77
Required tools and documentation	79
Guidelines for upgrading controllers	79
Learn about the ARL upgrade sequence	80
Stage 1. Prepare for upgrade	82
Stage 2. Relocate resources and retire node1	88
Stage 3. Boot node1 with the replacement system modules	109
Stage 4. Relocate resources and retire node2	121
Stage 5. Install the replacement system modules on node2	124
Stage 6. Boot node2 with the replacement system modules	134
Stage 7. Complete the upgrade	152
Troubleshoot	158
References	165
Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later	167
Learn about this ARL upgrade procedure	167
Automate the controller upgrade process	168
Decide whether to use this aggregate relocation procedure	169
Required tools and documentation	171

Guidelines for upgrading controllers with ARL	171
Verify the health of the MetroCluster configuration	173
Check for MetroCluster configuration errors	173
Verify switchover, healing, and switchback	174
Learn about the ARL upgrade sequence	174
Stage 1. Prepare for upgrade	176
Stage 2. Relocate and retire node1	182
Stage 3. Install and boot node3	185
Stage 4. Relocate and retire node2	212
Stage 5. Install and boot node4	214
Stage 6. Complete the upgrade	242
Troubleshoot	248
References	255
Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.5 to 9.7	257
Learn about this ARL upgrade procedure	257
Automate the controller upgrade process	258
Decide whether to use this aggregate relocation procedure	258
Required tools and documentation	260
Guidelines for upgrading controllers with ARL	260
Verify the health of the MetroCluster configuration	261
Check for MetroCluster configuration errors	262
Verify switchover, healing, and switchback	263
Learn about the ARL upgrade sequence	263
Stage 1. Prepare for upgrade	265
Stage 2. Relocate and retire node1	270
Stage 3. Install and boot node3	274
Stage 4. Relocate and retire node2	299
Stage 5. Install and boot node4	302
Stage 6. Complete the upgrade	326
Troubleshoot	333
References	340
Manually upgrade controller hardware running ONTAP 9.8 or later	342
Learn about this ARL upgrade procedure	342
Decide whether to use this aggregate relocation procedure	343
ARL upgrade workflow	344
Guidelines for upgrading controllers with ARL	347
Required tools and documentation	350
Worksheet: Information to collect before and during controller upgrade	350
Stage 1. Prepare for upgrade	352
Stage 2. Relocate and retire node1	373
Stage 3. Install and boot node3	387
Stage 4. Record information and retire node2	420
Stage 5. Install and boot node4	425
Stage 6. Complete the upgrade	452
Troubleshoot	457

References	464
Manually upgrade controller hardware running ONTAP 9.7 or earlier	466
Learn about this ARL upgrade procedure	466
Decide whether to use this aggregate relocation procedure	467
ARL upgrade workflow	468
Guidelines for upgrading controllers with ARL	471
Required tools and documentation	474
Worksheet: Information to collect before and during controller upgrade	474
Reconfigure the FC switch layout for ONTAP 9.1 or later	476
Stage 1. Prepare for upgrade	481
Stage 2. Relocate and retire node1	502
Stage 3. Install and boot node3	517
Stage 4. Record information and retire node2	552
Stage 5. Install and boot node4	557
Stage 6. Complete the upgrade	585
Troubleshoot	591
References	597

Upgrade by using ARL

Start here: Choose your ARL upgrade procedure

You can upgrade controller hardware without disruption by using aggregate relocation (ARL). For other methods of upgrading your controller hardware, see [Upgrade by moving volumes or storage](#).

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

ARL takes advantage of the HA configuration to give you the capability to move ownership of non-root aggregates from one node to another if they share storage within the same cluster.

There are several ARL methods for upgrading controller hardware. To select the appropriate procedure, review the following information on the systems and ONTAP versions supported for each ARL upgrade option.

Use "system controller replace" commands with ONTAP 9.15.1 or later

If your upgrade scenario is listed in the following supported systems matrix, go to [Use "system controller replace" commands to upgrade controller hardware introduced in ONTAP 9.15.1 and later](#) to start the upgrade procedure.

Show supported systems

Existing controller	Replacement controller	Supported beginning with ONTAP...
AFF A400	AFF A50	9.16.1
AFF A300	AFF A50	9.16.1
AFF A220, AFF A150	AFF A20	9.16.1
AFF C400	AFF C60, AFF C80	9.16.1
FAS8200, FAS8300, FAS9000	FAS70, FAS90, FAS50	9.15.1P3 for FAS70, FAS90 9.16.1P2 for FAS50
FAS8700	FAS70, FAS90	9.15.1P3
FAS9500	FAS90	9.15.1P3
AFF A300, AFF A400, AFF A700	AFF A70, AFF A90, AFF A1K	9.15.1
AFF A900	AFF A90, AFF A1K	9.15.1

Upgrade controller models in the same chassis

If your upgrade scenario is listed in the following supported systems matrix, go to [Use "system controller replace" commands to upgrade controller models in the same chassis](#) to start the upgrade procedure.

Show supported systems

Old system	Replacement system	Supported ONTAP versions
AFF C250	AFF C30, AFF C60	9.16.1 and later
AFF A250	AFF A50, AFF A30	9.16.1 and later
AFF C800	AFF C80	9.16.1 and later
AFF A800	AFF A70 or AFF A90	9.15.1 and later
AFF A220 configured as an All SAN Array (ASA)	ASA A150	9.13.1P1 and later
AFF A220	AFF A150	9.10.1P15, 9.11.1P11, 9.12.1P5 and later
AFF A200	AFF A150	9.10.1P15, 9.11.1P11 and later Note: AFF A200 does not support ONTAP versions later than 9.11.1.
AFF C190	AFF A150	9.10.1P15, 9.11.1P11, 9.12.1P5 and later
FAS2620	FAS2820	9.11.1P7 or later patch releases (FAS2620) Note: FAS2620 does not support ONTAP versions later than 9.11.1. 9.13.1 and later (FAS2820)
FAS2720	FAS2820	9.13.1 and later
AFF A700 configured as an ASA	ASA A900	9.13.1P1 and later
AFF A700	AFF A900	9.10.1P10, 9.11.1P6 and later
FAS9000	FAS9500	9.10.1P10, 9.11.1P6 and later

Use "system controller replace" commands with ONTAP 9.8 or later

If your upgrade scenario is listed in the following supported systems matrix, go to [Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later](#) to start the upgrade procedure.

Show supported systems

Old controller	Replacement controller
FAS8020, FAS8040, FAS8060, FAS8080	FAS8200, FAS8300, FAS8700, FAS9000
FAS8060, FAS8080	FAS9500
AFF8020, AFF8040, AFF8060, AFF8080	AFF A300, AFF A400, AFF A700, AFF A800
AFF8060, AFF8080	AFF A900
FAS8200	FAS8300, FAS8700, FAS9000, FAS9500
FAS8300, FAS8700, FAS9000	FAS9500
AFF A300	AFF A400, AFF A700, AFF A800, AFF A900
AFF A320	AFF A400
AFF A400, AFF A700	AFF A900

Use "system controller replace" commands with ONTAP 9.5 to 9.7

If your upgrade scenario is listed in the following supported systems matrix, go to [Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.5 to 9.7](#) to start the upgrade procedure.

Show supported systems

Old controller	Replacement controller
FAS8020, FAS8040, FAS8060, FAS8080	FAS8200, FAS8300, FAS8700, FAS9000
AFF8020, AFF8040, AFF8060, AFF8080	AFF A300, AFF A400, AFF A700, AFF A800
FAS8200	FAS8700, FAS9000, FAS8300
AFF A300	AFF A700, AFF A800, AFF A400

Use manual ARL commands

If your upgrade scenario is not supported using "system commands", you can perform an ARL upgrade using manual commands.

Show supported systems

ONTAP 9.8 or later

Manual ARL upgrades are supported for the following systems running ONTAP 9.8 and later:

- FAS system to FAS system
- AFF system to AFF system

You can only upgrade to a replacement system in the same series:

- AFF A-Series system to AFF A-Series system
- AFF C-Series system to AFF C-Series system
- ASA system to ASA system



ASA upgrades to an ASA r2 replacement system aren't supported. For information on migrating data from ASA to ASA r2, see [Enable data access from SAN hosts to your ASA r2 storage system](#).

You can only upgrade to a replacement system in the same series:

- ASA A-Series system to ASA A-Series system
- ASA C-Series system to ASA C-Series system

[Manually upgrade controller hardware running ONTAP 9.8 or later](#)

ONTAP 9.7 or earlier

Manual ARL upgrades are supported for the following systems running ONTAP 9.7 and earlier:

- FAS system to FAS system
- AFF system to AFF system

[Manually upgrade controller hardware running ONTAP 9.7 or earlier](#)

Use "system controller replace" commands to upgrade controller hardware introduced in ONTAP 9.15.1 and later

Learn about this ARL upgrade procedure

There are several aggregate relocation (ARL) methods for upgrading controller hardware. This ARL procedure describes how to upgrade an HA pair of controllers in AFF and FAS storage systems introduced in ONTAP 9.15.1 or later with new controllers while keeping the existing data and disks.



You cannot use this procedure to upgrade a MetroCluster FC or IP configuration. To upgrade a MetroCluster configuration, see [References](#) to link to the *MetroCluster Upgrade and Expansion documentation*.

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 confirm 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.

Terminology used in this information

In this information, the original nodes are called "node1" and "node2", and the new nodes are called "node3" and "node4". During the described procedure, node1 is replaced by node3, and node2 is replaced by node4.

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

Important information:

- This procedure is complex and assumes that you have advanced ONTAP administration skills. You also must read and understand the [guidelines for upgrading controllers with ARL](#) and the [ARL upgrade sequence](#) before beginning the upgrade.
- 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.
- 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 high-availability (HA) pair in the cluster.
- When you upgrade to a system introduced in ONTAP 9.15.1 or later, ONTAP converts the storage efficiency of existing volumes and applies the new storage efficiency features that make use of the hardware offload functionality. This is an automatic background process, with no visible performance impact to the system.
 - For AFF A20, AFF A30, AFF A50, AFF A70, AFF A90, AFF A1K, AFF C30, AFF C60, and AFF C80 systems, ONTAP converts the storage efficiency of all existing thin-provisioned volumes, including those not using storage efficiency.
 - For a FAS70 and FAS90 system, ONTAP only converts the storage efficiency of existing thin-provisioned volumes that had storage efficiency enabled before upgrading.

[Learn more about storage efficiency.](#)

- The AFF A20, AFF A50, AFF A70, AFF A90, AFF A1K, FAS70, and FAS90 systems share 100GbE network ports for both cluster and HA connections. These systems can support 10GbE or 25GbE cluster connections to legacy cluster switches; however, NetApp recommends updating to 100GbE cluster speeds when the 10GbE and 25GbE switches are no longer required. For more information, see the following Knowledge Base articles:
 - [How to configure 10G or 25G cluster ports on a new cluster setup](#)
 - [How to convert existing 10G or 25G cluster ports to 40G or 100G cluster ports](#)

The AFF A20 system shares 10GbE and 25GbE network ports for both cluster and HA connections. These are the only network port cluster connections supported by an AFF A20 system.

If you cannot link up e0a or e0b cluster ports on the existing node to the cluster ports on the new node, see [NetApp Bugs Online Bug ID CONTAP-166978](#).

Automate the controller upgrade process

During a controller upgrade, the controller is replaced with another controller running a newer or more powerful platform. This content provides the steps for the partially automated procedure, which utilizes automatic network port reachability checks to further simplify the controller upgrade experience.

Decide whether to use this aggregate relocation procedure


There are several aggregate relocation (ARL) methods for upgrading controller hardware. This ARL procedure describes how to upgrade an HA pair of controllers in AFF and FAS storage systems introduced in ONTAP 9.15.1 or later with new controllers while keeping the existing data and disks. This is a complex procedure that should be used only by experienced ONTAP administrators.

To help you decide if this ARL procedure is suitable for your controller hardware upgrade, you should review all of the following circumstances for supported and unsupported upgrades.

Supported upgrades for this ARL procedure

You can use this ARL procedure to upgrade a pair of nodes under the following circumstances:

- You're running ONTAP 9.15.1 or later.
- You don't want to add the new controllers as a new HA pair to the cluster and migrate the data using the volume move procedure.
- You're experienced in administering ONTAP and are comfortable with the risks of working in the diagnostic privilege mode.
- Your hardware upgrade combination is listed in the [supported model matrix](#).



You can use NetApp Storage Encryption (NSE), NetApp Volume Encryption (NVE), and NetApp Aggregate Encryption (NAE) with this procedure.

Upgrades not supported for this ARL procedure

You cannot use this ARL procedure to upgrade a pair of nodes under the following circumstances:

- You're performing one of the following upgrades:

Existing controller	Replacement controller
AFF A250	AFF A50, AFF A30
AFF A800	AFF A70, AFF A90
AFF C250	AFF C30, AFF C60
AFF C800	AFF C80

To perform upgrade an upgrade listed in the above table, see [References](#) to link to *Use "system controller replace" commands to upgrade controller models in the same chassis*.

- You're upgrading a MetroCluster FC or IP configuration. To upgrade a MetroCluster configuration, see [References](#) to link to the *MetroCluster Upgrade and Expansion documentation*.

Supported system upgrade combinations

The following table shows the supported controller upgrade combinations using ARL for AFF and FAS systems introduced in ONTAP 9.15.1 or later. If your controller upgrade combination isn't listed, contact technical support.

Existing controller	Replacement controller	Supported beginning with ONTAP...
AFF A400	AFF A50 ²	9.16.1
AFF A300	AFF A50 ²	9.16.1
AFF A220, AFF A150	AFF A20 ²	9.16.1 ¹
AFF C400	AFF C60, AFF C80	9.16.1
FAS8200, FAS8300, FAS9000	FAS70, FAS90, FAS50	9.15.1P3 for FAS70, FAS90 9.16.1P2 for FAS50
FAS8700	FAS70, FAS90	9.15.1P3
FAS9500	FAS90	9.15.1P3
AFF A300, AFF A400, AFF A700	AFF A70 ² , AFF A90 ² , AFF A1K	9.15.1
AFF A900	AFF A90 ² , AFF A1K	9.15.1

¹ To upgrade to an AFF A20, you must first convert an AFF A150 or AFF A220 to a DS224C shelf by swapping controller module with an IOM12 module. See [Convert an original node to a drive shelf](#).

² The AFF A20, AFF A50, AFF A70, and AFF A90 are integrated systems with two controllers in an HA configuration and onboard disks in a single chassis:

- If you're upgrading to an AFF A20, AFF A50, AFF A70, or AFF A90 with internal drives, you must remove ownership of these internal drives before you upgrade. After completing the upgrade, you can assign the internal drives to the AFF A20, AFF A50, AFF A70, or AFF A90 nodes and use them for creating data aggregates. You don't have to migrate root or data aggregates to internal drives.
- If you're upgrading to an AFF A20, AFF A50, AFF A70, or AFF A90 without internal drives, you don't need to assign internal drives after completing the upgrade.

Choose a different hardware upgrade procedure

- [Review the alternative ARL methods available for upgrading controller hardware](#).
- If you prefer a different method of upgrading the controller hardware and are willing to do volume moves, refer to [References](#) to link to *Upgrade by moving volumes or storage*.

Related information

Refer to [References](#) to link to the *ONTAP 9 Documentation*.

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 need the following tools to perform the upgrade:

- Grounding strap
- #2 Phillips screwdriver

Go to the [References](#) section to access the list of reference documents and reference sites required for this upgrade

Guidelines for upgrading controllers with ARL

To understand whether you can use ARL to upgrade a pair of controllers running ONTAP 9.15.1 or later depends on the platform and the configuration of both the original and replacement controllers.

Supported upgrades for ARL

Before upgrading a pair of nodes using this ARL procedure, review the following requirements to ensure your configuration is supported:

- Verify that ARL can be performed on the original and replacement controllers.
- Check the size of all defined aggregates and number of disks supported by the original system. You then compare the aggregate sizes and number of disks supported to the aggregate size and number of disks supported by the new system. Refer to [References](#) to link to the *Hardware Universe* where this information is available. 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.
- 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. For more information about cluster mixing rules, refer to [References](#) to link to the *Hardware Universe*.
- Migrate and re-home the cluster LIFs to two cluster ports per node if you have a system, such as an AFF 700, with the following configuration:
- More than two cluster ports per node
- A cluster interconnect card in slot4 in breakout mode to create ports e4a, e4b, e4c, and e4d, and ports e4e, e4f, e4g, and e4h



Performing a controller upgrade with more than two cluster ports per node might result in missing cluster LIFs on the new controller after the upgrade.

For more information, see the Knowledge Base article [How to delete unwanted or unnecessary cluster LIFs](#).

Controller upgrade using ARL is supported on systems configured with SnapLock Enterprise and SnapLock Compliance volumes.

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.

Upgrades not supported for ARL

You cannot upgrade replacement controllers that do not support the disk shelves connected to the original controllers.

Refer to [References](#) to link to the *Hardware Universe* for disk-support information.

If you want to upgrade entry level controllers with internal drives, refer to [References](#) to link to *Upgrade by moving volumes or storage* and go to the procedure *Upgrading a pair of nodes running clustered Data ONTAP by moving volumes*.

Troubleshoot

If any problems occur while upgrading the controllers, see [Troubleshoot](#) for more information and possible solutions.

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

Learn about the ARL upgrade sequence

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

Upgrade 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.

ARL upgrade sequence overview

During the procedure, you upgrade the original controller hardware with the replacement controller hardware, one controller at a time, taking advantage of the HA pair 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 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:

Stage	Description
Stage 1. Prepare for upgrade	<p>During Stage 1, you run prechecks and, if required, correct aggregate ownership. You must record certain information if you are managing storage encryption by using the OKM and you can choose to quiesce the SnapMirror relationships.</p> <p>Aggregate ownership at the end of Stage 1:</p> <ul style="list-style-type: none"> • Node1 is the home owner and current owner of the node1 aggregates. • Node2 is the home owner and current owner of the node2 aggregates.
Stage 2. Relocate and retire node1	<p>During Stage 2, you relocate node1 non-root aggregates and NAS data LIFs to node2. This process is largely automated; the operation pauses to enable you to check its status. You must manually resume the operation. If required, you relocate failed or vetoed aggregates. You record node1 information for use later in the procedure before retiring node1. You can also prepare to netboot node3 and node4 later in the procedure.</p> <p>Aggregate ownership at the end of Stage 2:</p> <ul style="list-style-type: none"> • Node2 is the current owner of node1 aggregates. • Node2 is the home owner and current owner of node2 aggregates.
Stage 3. Install and boot node3	<p>During Stage 3, you install and boot node3, check that the cluster and node-management ports from node1 come online on node3, and verify the node3 installation. If you are using NetApp Volume Encryption (NVE), you restore key-manager configuration. You also relocate the node1 NAS data LIFs and non-root aggregates from node2 to node3 and verify that the SAN LIFs exist on node3.</p> <p>Aggregate ownership at the end of Stage 3:</p> <ul style="list-style-type: none"> • Node3 is the home owner and current owner of node1 aggregates. • Node2 is the home owner and current owner of node2 aggregates.
Stage 4. Relocate and retire node2	<p>During Stage 4, you relocate non-root aggregates and NAS data LIFs from node2 to node3. You also record node2 information for use later in the procedure before retiring node2.</p> <p>Aggregate ownership at the end of Stage 4:</p> <ul style="list-style-type: none"> • Node3 is the home owner and current owner of aggregates that originally belonged to node1. • Node2 is the home owner of node2 aggregates. • Node3 is the current owner of node2 aggregates.

Stage	Description
Stage 5. Install and boot node4	<p>During Stage 5, you install and boot node4, check that the cluster and node-management ports from node2 come online on node4, and verify the node4 installation. If you are using NVE, you restore key-manager configuration. You also relocate node2 NAS data LIFs and non-root aggregates from node3 to node4 and verify that the SAN LIFs exist on node4.</p> <p>Aggregate ownership at the end of Stage 5:</p> <ul style="list-style-type: none"> • Node3 is the home owner and current owner of the aggregates that originally belonged to node1. • Node4 is the home owner and current owner of aggregates that originally belonged to node2.
Stage 6. Complete the upgrade	<p>During Stage 6, you confirm that the new nodes are set up correctly and, if the new nodes are encryption-enabled, you configure and set up Storage Encryption or NVE. You should also decommission the old nodes and resume the SnapMirror operations.</p>

Stage 1. Prepare for upgrade

Prepare the nodes for upgrade

The controller replacement process begins with a series of prechecks. You also gather information about the original nodes for use later in the procedure and, if required, determine the type of self-encrypting drives that are in use.

Steps

1. Begin the controller replacement process by entering the following command in the ONTAP command line:

```
system controller replace start -nodes <node_names>
```



You can only execute the system controller replace start command at the advanced privilege level: `set -privilege advanced`

You will see output similar to the following example. The output displays the ONTAP version running on your cluster:

Warning: 1. Current ONTAP version is 9.15.1

2. Verify that NVMEM or NVRAM batteries of the new nodes are charged, and charge them if they are not. You need to physically check the new nodes to see if the NVMEM or NVRAM batteries are charged. You can check the battery status either by connecting to a serial console or using SSH, logging into the Service Processor (SP) or Baseboard Management Controller (BMC) 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.

3. If a controller was previously part of a different cluster, run wipeconfig before using it as the replacement controller.

4. Note: This is not a MetroCluster configuration. Controller replacement supports only ARL based procedure.
Do you want to continue? {y|n}: y

2. Press `y`, you will see the following output:

```
Controller replacement operation: Prechecks in progress.  
Controller replacement operation has been paused for user intervention.
```

The system runs the following prechecks; record the output of each precheck for use later in the procedure:

Precheck	Description
Cluster Health Check	Checks all the nodes in the cluster to confirm they are healthy.
Aggregate Relocation Status Check	Checks whether an aggregate relocation is already in progress. If another aggregate relocation is in progress, the check fails.
Model Name Check	Checks whether the controller models are supported for this procedure. If the models are not supported, the task fails.
Cluster Quorum Check	Checks that the nodes being replaced are in quorum. If the nodes are not in quorum, the task fails.
Image Version Check	Checks that the nodes being replaced run the same version of ONTAP. If the ONTAP image versions are different, the task fails. 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. For instructions on how to upgrade ONTAP, refer to References to link to <i>Upgrade ONTAP</i> .

Precheck	Description
HA Status Check	Checks if both the nodes being replaced are in a high- availability (HA) pair configuration. If storage failover is not enabled for the controllers, the task fails.
Aggregate Status Check	If the nodes being replaced own aggregates for which they are not the home owner, the task fails. The nodes should not own any non-local aggregates.
Disk Status Check	If any nodes being replaced have missing or failed disks, the task fails. If any disks are missing, refer to References to link to <i>Disk and aggregate management with the CLI</i> , <i>Logical storage management with the CLI</i> , and <i>HA pair management</i> to configure storage for the HA pair.
Data LIF Status Check	Checks if any of the nodes being replaced have non- local data LIFs. The nodes should not contain any data LIFs for which they are not the home owner. If one of the nodes contains non-local data LIFs, the task fails.
Cluster LIF Status	Checks whether the cluster LIFs are up for both nodes. If the cluster LIFs are down, the task fails.
ASUP Status Check	If ASUP notifications are not configured, the task fails. You must enable ASUP before beginning the controller replacement procedure.
CPU Utilization Check	Checks if the CPU utilization is more than 50% for any of the nodes being replaced. If the CPU usage is more than 50% for a considerable period of time, the task fails.
Aggregate Reconstruction Check	Checks if reconstruction is occurring on any data aggregates. If aggregate reconstruction is in progress, the task fails.
Node Affinity Job Check	Checks if any node affinity jobs are running. If node affinity jobs are running, the check fails.

- After the controller replacement operation is started and the prechecks are completed, the operation pauses enabling you to collect output information that you might need later when configuring node3.



Before you start the upgrade, you migrate and re-home the cluster LIFs to two cluster ports per node if you have a system, such as an AFF 700, with the following configuration:

- More than two cluster ports per node
- A cluster interconnect card in slot4 in breakout mode to create ports e4a, e4b, e4c, and e4d, and ports e4e, e4f, e4g, and e4h

Performing a controller upgrade with more than two cluster ports per node might result in missing cluster LIFs on the new controller after the upgrade.

For more information, see the Knowledge Base article [How to delete unwanted or unnecessary cluster LIFs](#).

4. Run the below set of commands as directed by the controller replacement procedure on the system console.

From the serial port connected to each node, run and save the output of the following commands individually:

```
° vservers services name-service dns show
° network interface show -curr-node <local> -role <cluster,intercluster,node-
  mgmt,cluster-mgmt,data>
° network port show -node <local> -type physical
° service-processor show -node <local> -instance
° network fcp adapter show -node <local>
° network port ifgrp show -node <local>
° system node show -instance -node <local>
° run -node <local> sysconfig
° storage aggregate show -r
° storage aggregate show -node <local>
° volume show -node <local>
° system license show -owner <local>
° storage encryption disk show
° security key-manager onboard show-backup
° security key-manager external show
° security key-manager external show-status
° network port reachability show -detail -node <local>
```



If NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE) using the Onboard Key Manager (OKM) is in use, keep the key manager passphrase ready to complete the key manager resync later in the procedure.

5. If your system uses self-encrypting drives, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use on the HA pair that you are upgrading. ONTAP software supports two types of self-encrypting drives:
 - ° FIPS-certified NetApp Storage Encryption (NSE) SAS or NVMe drives
 - ° Non-FIPS self-encrypting NVMe drives (SED)

[Learn more about supported self-encrypting drives.](#)

Correct aggregate ownership if an ARL precheck fails

If the Aggregate Status Check fails, you must return aggregates owned by the partner node to the home owner node and initiate the precheck process again.

Steps

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

```
storage aggregate relocation start -node source_node -destination destination-  
node -aggregate-list *
```

2. 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
```

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.
```

After you finish

You must restart the controller replacement process:

```
system controller replace start -nodes node_names
```

License

For detailed information about ONTAP licensing, refer to [License management](#).



Using unlicensed features on the controller might put you out of compliance with your license agreement.

Manage storage encryption using the Onboard Key Manager

You can use the Onboard Key Manager (OKM) to manage encryption keys. If you have the OKM set up, you must record the passphrase and backup material before beginning the upgrade.

Steps

1. Record the cluster-wide passphrase.

This is the passphrase that was entered when the OKM was configured or updated using the CLI or REST API.

2. Back up the key-manager information by running the `security key-manager onboard show-backup` command.

Quiesce the SnapMirror relationships (optional)

Before continuing with the procedure, you must confirm that all the SnapMirror relationships are quiesced. When a SnapMirror relationship is quiesced, it remains quiesced across reboots and failovers.

Steps

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

```
snapmirror show
```



If the status is "Transferring", you must abort those transfers:

```
snapmirror abort -destination-vserver vservice_name
```

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

2. Quiesce all relationships between the cluster:

```
snapmirror quiesce -destination-vserver *
```

Stage 2. Relocate and retire node1

Relocate non-root aggregates and NAS data LIFs owned by node1 to node2

Before you can replace node1 with node3, you must move the non-root aggregates and NAS data LIFs from node1 to node2 before eventually moving node1's resources to node3.

Before you begin

The operation should already be paused when you begin the task; you must manually resume the operation.

About this task

After the aggregates and LIFs are migrated, the operation is paused for verification purposes. At this stage, you must verify whether or not all the non-root aggregates and non-SAN data LIFs are migrated to node3.



The home owner for the aggregates and LIFs is not modified; only the current owner is modified.

Steps

1. Resume the aggregate relocation and NAS data LIF move operations:

```
system controller replace resume
```

All the non-root aggregates and NAS data LIFs are migrated from node1 to node2.

The operation pauses to enable you to verify whether all node1 non-root aggregates and non-SAN data LIFs have been migrated to node2.

2. Check the status of the aggregate relocation and NAS data LIF move operations:

```
system controller replace show-details
```

3. With the operation still paused, verify that all the non-root aggregates are online for their state on node2:

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

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,normal							
aggr_2	825.0GB	825.0GB	0%	online	1	node2	
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 <aggregate_name>
```

4. Verify that all the volumes are online on node2 by using 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 <vserver_name> -volume <volume_name>
```

The `vserver_name` to use with this command is found in the output of the previous `volume show` command.

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

```
network interface modify -vserver vserver_name -lif LIF_name -home-node  
nodename -status-admin up
```

Relocate failed or vetoed aggregates to node2

If any aggregates fail to relocate or are vetoed, you need to manually relocate the aggregates to node2, or if necessary, override either the vetoes or destination checks.

About this task

The system pauses the relocation operation due to the error.

Steps

1. Check the event management system (EMS) logs to determine why the aggregate failed to relocate or was vetoed.
2. Relocate any failed or vetoed aggregates:

```
storage aggregate relocation start -node <node1> -destination <node2>  
-aggregate-list <aggregate_name> -ndo-controller-upgrade true
```

3. When prompted, enter `y`.
4. You can force relocation by using one of the following methods:

Option	Description
Overriding veto checks	Use the following command: storage aggregate relocation start -node node1 -destination node2 -aggregate-list <aggregate_list> -ndo-controller-upgrade true -override-vetoes true
Overriding destination checks	Use the following command: storage aggregate relocation start -node node1 -destination node2 -aggregate-list <aggregate_list> -ndo-controller-upgrade true -override-vetoes true -override-destination-checks true

Retire node1

To retire node1, you resume the automated operation to disable the HA pair with node2 and shut down node1 correctly. Later in the procedure, you remove node1 from the rack or chassis.

Steps

1. Resume the operation:

```
system controller replace resume
```

2. Verify that node1 has been halted:

```
system controller replace show-details
```

After you finish

After retiring node1, verify that it's powered off and not connected to the network. Leave the old hardware of node1 in this state until you complete the upgrade of both node1 and node2. You can then decommission node1 during [Stage 6. Complete the upgrade](#).

Prepare for netboot

After you physically rack node3 and node4 later in the procedure, you might need to netboot them. The term "netboot" means you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you put a copy of the ONTAP 9 boot image onto a web server that the system can access.

You can also use the USB boot option to perform a netboot. See the Knowledge Base article [How to use the boot_recovery LOADER command for installing ONTAP for initial setup of a system](#).

Before you begin

- Verify that you can access a HTTP server with the system.
- Refer to [References](#) to link to the *NetApp Support Site* and download the necessary system files for your platform and the correct version of ONTAP.

About this task

You must netboot the new controllers if they do not have the same 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.

Steps

1. Access the NetApp Support Site to download the files used for performing the netboot of the system.
2. Download the appropriate ONTAP software from the software download section of the NetApp Support Site and store the `<ontap_version>_image.tgz` file on a web-accessible directory.
3. Change to the web-accessible directory and verify that the files you need are available.

Your directory listing should contain the following file:

`<ontap_version>_image.tgz`



You do not need to extract the contents of the `<ontap_version>_image.tgz` file.

You will use the information in the directories in [Stage 3](#).

Stage 3. Install and boot node3

Install and boot node3

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

About this task

The relocation operation is paused at the beginning of this stage. This process is largely automated; the operation pauses to enable you to check its status. You must manually resume the operation. In addition, you must verify that the SAN LIFs have successfully come online and are assigned to the correct FC physical ports on node3.

You need to netboot node3 if it does not have the same version of ONTAP 9 that is installed on node1. After you install node3, boot it 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, by following the instructions in [Prepare for netboot](#).

Steps

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

The space and height requirements for the new nodes might be different from the existing nodes. Plan for the space requirements for your upgrade scenario.

2. Install node3 in the rack, following the *Installation and Setup Instructions* for your node model.
3. Cable node3, moving the connections from node1 to node3.

Beginning with ONTAP 9.15.1, new controller models have only one "wrench" port for the baseboard management controller (BMC) and management connections. Plan the cabling changes accordingly.

- Console (remote management port)
- Cluster and HA ports
- Data ports
- Cluster and node management ports
- Serial-attached SCSI (SAS) and Ethernet storage ports
- SAN configurations: iSCSI Ethernet, FC, and NVMe/FC switch ports



You might need to change the interconnect cables between the old and new controllers to allow interoperability between the different controller and card models. Refer to the [system installation procedures](#) for a cabling map of the Ethernet storage shelves for your systems.

For controllers introduced in ONTAP 9.15.1 and later, cluster and HA interconnects use the same ports. For switch connected configurations, it is required to connect similar ports to the same cluster switches. For example, when upgrading to an AFF A1K from an existing controller, you should connect e1a ports on both nodes to one switch and e7a ports on both nodes to the second switch.

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.



When you boot node3, you might see the following warning 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.



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




Refer to [Prepare for netboot](#).

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



You must 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 using the following command at the boot environment prompt: <code>ifconfig e0M -auto</code>
Not running	<p>Manually configure the connection by using the following command at the boot environment prompt:</p> <pre>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></pre> <p><i>filer_addr</i> is the IP address of the storage system (mandatory). <i>netmask</i> is the network mask of the storage system (mandatory). <i>gateway</i> is the gateway for the storage system (mandatory). <i>dns_addr</i> is the IP address of a name server on your network (optional). <i>dns_domain</i> is the Domain Name Service (DNS) domain name (optional).</p> <div>  <p>Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p> </div>

7. Perform netboot on node3:

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

The `<path_to_the_web-accessible_directory>` should lead to where you downloaded the `<ontap_version>_image.tgz` in the section [Prepare for netboot](#).



Do not interrupt the boot.

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

This menu option downloads and installs the new 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 ONTAP, and not upgrades of controllers.



Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the incorrect image might install. This issue applies to all ONTAP releases. The netboot procedure combined with option (7) `Install new software` wipes the boot media and places the same ONTAP version on both image partitions.

9. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL:

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

10. Complete the following substeps to reboot the controller module:

- 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}
```

- b. Enter `y` to reboot 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 must be restored.

11. Select maintenance mode 5 from the boot menu and enter `y` when you are prompted to continue with the boot.
12. Verify that the controller and chassis are configured as ha:

```
ha-config show
```

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

```
Chassis HA configuration: ha  
Controller HA configuration: ha
```



System records 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.

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

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

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

14. Confirm that all Ethernet ports used to connect to the Ethernet shelves are configured as storage:

```
storage port show
```

The output displayed depends on the system configuration. The following output example is for a node with a single storage card in slot11. The output for your system might be different:

```
*> storage port show
Port Type Mode      Speed (Gb/s) State      Status  VLAN ID
---- ---- -
e11a ENET storage 100 Gb/s    enabled  online  30
e11b ENET storage 100 Gb/s    enabled  online  30
```

15. Modify the ports that are not set to storage:

```
storage port modify -p <port> -m storage
```

All Ethernet ports connected to storage shelves must be configured as storage to allow access to the disks and shelves.

16. Exit maintenance mode:

```
halt
```

Interrupt the autoboot by pressing **Ctrl-C** at the boot environment prompt.

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

```
date
```

18. On node3, check the date by using the following command at the boot environment prompt:

```
show date
```

19. If necessary, set the date on node3:

```
set date <mm/dd/yyyy>
```

20. On node3, check the time by using the following command at the boot environment prompt:

```
show time
```

21. If necessary, set the time on node3:

```
set time <hh:mm:ss>
```

22. In boot loader, set the partner system ID on node3:

```
setenv partner-sysid <node2_sysid>
```

For node3, partner-sysid must be that of node2.

- a. Save the settings:

```
saveenv
```

23. Verify the partner-sysid for node3:

```
printenv partner-sysid
```

24. If you have NetApp Storage Encryption (NSE) drives installed, perform the following steps.



If you have not already done so earlier in the procedure, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use.

- a. Set bootarg.storageencryption.support to true or false:

If the following drives are in use...	Then...
NSE drives that conform to FIPS 140-2 Level 2 self-encryption requirements	setenv bootarg.storageencryption.support true
NetApp non-FIPS SEDs	setenv bootarg.storageencryption.support false

- b. Go to the special boot menu and select option (10) Set Onboard Key Manager recovery secrets.

Enter the passphrase and the backup information that you recorded earlier procedure. See [Manage storage encryption using the Onboard Key Manager](#).

25. Boot the node into boot menu:

```
boot_ontap menu
```

26. When node3 stops at the boot menu, reassign the node1 disks to node3 by running the following command on node3:

```
boot_after_controller_replacement
```

After a short delay, you are prompted to enter the name of the node that is being replaced. If there are shared disks (also called Advanced Disk Partitioning (ADP) or partitioned disks), you are prompted to enter the node name of the HA partner.

These prompts might get buried in the console messages. If you do not enter a node name or enter an incorrect name, you are prompted to enter the name again.

Expand the console output example

```
LOADER-A> boot_ontap menu
.
<output truncated>
.
All rights reserved.
*****
*                                     *
* Press Ctrl-C for Boot Menu. *
*                                     *
*****
.
<output truncated>
.
Please choose one of the following:
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 22/7
(22/7) Print this secret List
(25/6) Force boot with multiple filesystem disks missing.
(25/7) Boot w/ disk labels forced to clean.
(29/7) Bypass media errors.
(44/4a) Zero disks if needed and create new flexible root volume.
(44/7) Assign all disks, Initialize all disks as SPARE, write DDR
labels
.
<output truncated>
.
(wipeconfig) Clean all configuration on boot
device
(boot_after_controller_replacement) Boot after controller upgrade
(boot_after_mcc_transition) Boot after MCC transition
(9a) Unpartition all disks and remove
their ownership information.
(9b) Clean configuration and
initialize node with partitioned disks.
```

```

(9c)                                Clean configuration and
initialize node with whole disks.
(9d)                                Reboot the node.
(9e)                                Return to main boot menu.
The boot device has changed. System configuration information could
be lost. Use option (6) to restore the system configuration, or
option (4) to initialize all disks and setup a new system.
Normal Boot is prohibited.
Please choose one of the following:
(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? boot_after_controller_replacement
This will replace all flash-based configuration with the last backup
to disks. Are you sure you want to continue?: yes
.
<output truncated>
.
Controller Replacement: Provide name of the node you would like to
replace:<nodename of the node being replaced>
Controller Replacement: Provide High Availability partner of node1:
<nodename of the partner of the node being replaced>
Changing sysid of node node1 disks.
Fetched sanown old_owner_sysid = 536940063 and calculated old sys id
= 536940063
Partner sysid = 4294967295, owner sysid = 536940063
.
<output truncated>
.
varfs_backup_restore: restore using /mroot/etc/varfs.tgz
varfs_backup_restore: attempting to restore /var/kmip to the boot
device
varfs_backup_restore: failed to restore /var/kmip to the boot device
varfs_backup_restore: attempting to restore env file to the boot
device
varfs_backup_restore: successfully restored env file to the boot
device wrote key file "/tmp/rndc.key"
varfs_backup_restore: timeout waiting for login

```

```

varfs_backup_restore: Rebooting to load the new varfs
Terminated
<node reboots>
System rebooting...
.
Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...
Rebooting to load the restored env file...
.
System rebooting...
.
<output truncated>
.
WARNING: System ID mismatch. This usually occurs when replacing a
boot device or NVRAM cards!
Override system ID? {y|n} y
.
Login:

```



In the above console output example, ONTAP will prompt you for the partner node name if the system uses Advanced Disk Partitioning (ADP) disks.

27. If the system goes into a reboot loop with the message `no disks found`, it indicates that there was a problem with the disk reassignment. See [Troubleshoot](#) to resolve the issue.
28. Press `Ctrl-C` during autoboot to stop the node at the `LOADER>` prompt.
29. At the `LOADER` prompt, enter maintenance mode:

```
boot_ontap maint
```

30. Verify the disk connectivity, controller model string, HA-configuration, and other hardware connectivity related details.
31. Exit maintenance mode:

```
halt
```

32. At the `LOADER` prompt, boot up:

```
boot_ontap menu
```

Now, on booting, the node can detect all the disks that were previously assigned to it and can boot up as expected.

When the cluster nodes you are replacing use root volume encryption, ONTAP is unable to read the volume information from the disks. Restore the keys for the root volume.



This only applies when the root volume is using NetApp Volume Encryption.

- a. Return to the special boot menu:

```
LOADER> boot_ontap menu
```

```
Please choose one of the following:
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.

Selection (1-11)? 10
```

- b. Select **(10) Set Onboard Key Manager recovery secrets**

- c. Enter **y** at the following prompt:

```
This option must be used only in disaster recovery procedures. Are you sure?
(y or n): y
```

- d. At the prompt, enter the key-manager passphrase.

- e. Enter the backup data when prompted.



You must have obtained the passphrase and backup data in the [Prepare the nodes for upgrade](#) section of this procedure.

- f. After the system boots to the special boot menu again, run option **(1) Normal Boot**



You might encounter an error at this stage. If an error occurs, repeat the substeps in [Step 32](#) until the system boots normally.

Verify the node3 installation

You must verify that the physical ports from node1 map correctly to the physical ports on node3. This will enable node3 to communicate with other nodes in the cluster and with the network after the upgrade.

About this task

Refer to [References](#) to link to the *Hardware Universe* to capture information about the ports on the new nodes. You will use the information later in this section.

Physical port layout might vary, depending on the model of the nodes. When the new node boots up, ONTAP will try to determine which ports should host cluster LIFs in order to automatically come into quorum.

If the physical ports on node1 do not map directly to the physical ports on node3, the subsequent section [Restore network configuration on node3](#) must be used to repair the network connectivity.

After you install and boot node3, you must verify that it is installed correctly. You must wait for node3 to join quorum and then resume the relocation operation.

At this point in the procedure, the operation will have paused as node3 joins quorum.

Steps

1. Verify that node3 has joined quorum:

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

The output of the `health` field should be `true`.

2. Verify that node3 is part of the same cluster as node2 and that it is healthy:

```
cluster show
```

3. Switch to advanced privilege mode:

```
set advanced
```

4. Check the status of the controller replacement operation and verify that it is in a paused state and in the same state that it was in before node1 was halted to perform the physical tasks of installing new controllers and moving cables:

```
system controller replace show
```

```
system controller replace show-details
```

5. Resume the controller replacement operation:

```
system controller replace resume
```

6. Controller replacement will pause for intervention with the following message:

```

Cluster::*> system controller replace show
Node              Status              Error-Action
-----
Node1(now node3) Paused-for-intervention Follow the instructions
given in
Step Details
Node2              None
Step Details:
-----
To complete the Network Reachability task, the ONTAP network
configuration must be manually adjusted to match the new physical
network configuration of the hardware. This includes:

1. Re-create the interface group, if needed, before restoring VLANs. For
detailed commands and instructions, refer to the "Re-creating VLANs,
ifgrps, and broadcast domains" section of the upgrade controller
hardware guide for the ONTAP version running on the new controllers.
2. Run the command "cluster controller-replacement network displaced-
vlans show" to check if any VLAN is displaced.
3. If any VLAN is displaced, run the command "cluster controller-
replacement network displaced-vlans restore" to restore the VLAN on the
desired port.

2 entries were displayed.

```



In this procedure, the section *Re-creating VLANs, ifgrps, and broadcast domains* has been renamed *Restore network configuration on node3*.

7. With the controller replacement in a paused state, proceed to the next section of this document to restore network configuration on the node.

Restore network configuration on node3

After you confirm that node3 is in quorum and can communicate with node2, verify that node1's VLANs, interface groups, and broadcast domains are seen on node3. Also, verify that all node3 network ports are configured in their correct broadcast domains.

About this task

For more information on creating and re-creating VLANs, interface groups, and broadcast domains, refer to [References](#) to link to *Network Management*.

Steps

1. List all the physical ports that are on upgraded node1 (referred to as node3):

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

All physical network ports, VLAN ports and interface group ports on the node are displayed. From this output, you can see any physical ports that have been moved into the `Cluster` broadcast domain by ONTAP. You can use this output to aid in deciding which ports must be used as interface group member ports, VLAN base ports or standalone physical ports for hosting LIFs.

2. List the broadcast domains on the cluster:

```
network port broadcast-domain show
```

3. List the network port reachability of all ports on node3:

```
network port reachability show
```

You should see output like the following example:

```
ClusterA::*> network port reachability show
Node      Port      Expected Reachability      Reachability
Status
-----
node1_node3
      e0M      Default:Mgmt      ok
      e10a      Default:Default      ok
      e10b      -      no-reachability
      e10c      Default:Default      ok
      e10d      -      no-reachability
      e1a      Cluster:Cluster      ok
      e1b      -      no-reachability
      e7a      Cluster:Cluster      ok
      e7b      -      no-reachability
node2_node4
      e0M      Default:Mgmt      ok
      e4a      Default:Default      ok
      e4b      -      no-reachability
      e4c      Default:Default      ok
      e4d      -      no-reachability
      e3a      Cluster:Cluster      ok
      e3b      Cluster:Cluster      ok
18 entries were displayed.
```

In the preceding example, `node1_node3` is just booted after controller replacement. Some ports do not have reachability to their expected broadcast domains and must be repaired.

4. Repair the reachability for each of the ports on node3 with a reachability status other than `ok`. Run the following command, first on any physical ports, then on any VLAN ports, one at a time:

```
network port reachability repair -node <node_name> -port <port_name>
```

You should see output like the following example:

```
Cluster ::> reachability repair -node node1_node3 -port e4a
```

```
Warning: Repairing port "node1_node3: e4a" may cause it to move into a
different broadcast domain, which can cause LIFs to be re-homed away
from the port. Are you sure you want to continue? {y|n}:
```

A warning message, as shown above, is expected for ports with a reachability status that might be different from the reachability status of the broadcast domain where it is currently located. Review the connectivity of the port and answer `y` or `n` as appropriate.

Verify that all physical ports have their expected reachability:

```
network port reachability show
```

As the reachability repair is performed, ONTAP attempts to place the ports in the correct broadcast domains. However, if a port's reachability cannot be determined and does not belong to any of the existing broadcast domains, ONTAP will create new broadcast domains for these ports.

5. If interface group configuration does not match the new controller physical port layout, modify it by using the following steps.
 - a. You must first remove physical ports that should be interface group member ports from their broadcast domain membership. You can do this by using the following command:

```
network port broadcast-domain remove-ports -broadcast-domain <broadcast-
domain_name> -ports <node_name:port_name>
```

- b. Add a member port to an interface group:

```
network port ifgrp add-port -node <node_name> -ifgrp <ifgrp> -port
<port_name>
```

- c. The interface group is automatically added to the broadcast domain about a minute after the first member port is added.
 - d. Verify that the interface group was added to the appropriate broadcast domain:

```
network port reachability show -node <node_name> -port <ifgrp>
```

If the interface group's reachability status is not `ok`, assign it to the appropriate broadcast domain:

```
network port broadcast-domain add-ports -broadcast-domain
<broadcast_domain_name> -ports <node:port>
```

6. Assign appropriate physical ports to the `Cluster` broadcast domain by using the following steps:
 - a. Determine which ports have reachability to the `Cluster` broadcast domain :

```
network port reachability show -reachable-broadcast-domains Cluster:Cluster
```

- b. Repair any port with reachability to the `Cluster` broadcast domain, if its reachability status is not `ok`:

```
network port reachability repair -node <node_name> -port <port_name>
```

7. Move the remaining physical ports into their correct broadcast domains by using one of the following commands:

```
network port reachability repair -node <node_name> -port <port_name>
```

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

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

Verify that there are no unreachable or unexpected ports present. Check the reachability status for all physical ports by using the following command and examining the output to confirm the status is `ok`:

```
network port reachability show -detail
```

8. Restore any VLANs that might have become displaced by using the following steps:

- a. List displaced VLANs:

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

Output like the following should display:

```
Cluster::*> displaced-vlans show
(cluster controller-replacement network displaced-vlans show)
      Original
Node   Base Port   VLANs
-----
Node1   a0a           822, 823
        e4a           822, 823
2 entries were displayed.
```

- b. Restore VLANs that were displaced from their previous base ports:

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

The following is an example of restoring VLANs that have been displaced from interface group "a0a" back onto the same interface group:

```
Cluster::*> displaced-vlans restore -node node1_node3 -port a0a
-destination-port a0a
```

The following is an example of restoring displaced VLANs on port "e9a" to 'e9d':

```
Cluster::*> displaced-vlans restore -node node1_node3 -port e9a
-destination-port e9d
```

When a VLAN restore is successful, the displaced VLANs are created on the specified destination port. The VLAN restore fails if the destination port is a member of an interface group, or if the destination port is down.

Wait about one minute for newly restored VLANs to be placed into their appropriate broadcast domains.

- c. Create new VLAN ports as needed for VLAN ports that are not in the `cluster controller-replacement network displaced-vlans show` output but should be configured on other physical ports.

9. Delete any empty broadcast domains after all port repairs have been completed:

```
network port broadcast-domain delete -broadcast-domain <broadcast_domain_name>
```

10. Verify port reachability:

```
network port reachability show
```

When all ports are correctly configured and added to the correct broadcast domains, the `network port reachability show` command should report the reachability status as `ok` for all connected ports, and the status as `no-reachability` for ports with no physical connectivity. If any port reports a status other than these two, perform the reachability repair and add or remove ports from their broadcast domains as instructed in [Step 4](#).

11. Verify that all ports have been placed into broadcast domains:

```
network port show
```

12. Verify that all ports in the broadcast domains have the correct maximum transmission unit (MTU) configured:

```
network port broadcast-domain show
```

13. Restore LIF home ports, specifying the Vserver(s) and LIF(s) home ports, if any, that need to be restored by using the following steps:

- a. List any LIFs that are displaced:

```
displaced-interface show
```

- b. Restore LIF home nodes and home ports:

```
cluster controller-replacement network displaced-interface restore-home-node
-node <node_name> -vserver <vserver_name> -lif-name <LIF_name>
```

14. Verify that all LIFs have a home port and are administratively up:

```
network interface show -fields home-port, status-admin
```

Restore key-manager configuration on node3

If you are using NetApp Volume Encryption (NVE) and NetApp Aggregate Encryption (NAE) to encrypt volumes on the system you are upgrading, the encryption configuration must be synchronized to the new nodes. If you do not synchronize the key-manager, when you relocate the node1 aggregates from node2 to node3 by using ARL, failures might occur because node3 does not have the required encryption keys to bring encrypted volumes and aggregates online.

About this task

Synchronize the encryption configuration to the new nodes by performing the following steps:

Steps

- 1. Run the following command from node3:
- 2. Verify that the SVM-KEK key is restored to "true" on node3 before you relocate the data aggregates:

```
::> security key-manager onboard sync
```

```
::> security key-manager key query -node node3 -fields restored -key -type SVM-KEK
```

Example

```
::> security key-manager key query -node node3 -fields restored -key -type SVM-KEK
```

node	vserver	key-server	key-id
restored			
-----	-----	-----	-----
node3	svm1	""	0000000000000000020000000000a008a81976
true			2190178f9350e071fbb90f000000000000000

Move non-root aggregates and NAS data LIFs owned by node1 from node2 to node3

After you verify network configuration on node3 and before you relocate aggregates from node2 to node3, you must verify that the NAS data LIFs belonging to node1 that are currently on node2 are relocated from node2 to node3. You must also verify that the SAN LIFs exist 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. The iSCSI LIFs automatically find the correct home ports using the reachability scan. The FC and NVMe/FC SAN LIFs do not move automatically. They continue to show the home port they were on before upgrading.

Check the SAN LIFs on node3:

- a. Modify any iSCSI SAN LIFs reporting a "down" operation status to the new data ports:

```
network interface modify -vserver <vserver> -lif <iscsi_san_lif>  
admin down
```

```
network interface modify -vserver <vserver> -lif <iscsi_san_lif> port  
<new_port> node <node>
```

```
network interface modify -vserver <vserver> -lif <iscsi_san_lif>
```

- b. Modify any FC and NVMe/FC SAN LIFs that are home to the new controller and reporting a "down" operational status to the FCP ports on the new controller:

```
network interface modify -vserver <vserver> -lif <fc_san_lif> admin  
down
```

```
network interface modify -vserver <vserver> -lif <fc_san_lif> port  
<new_port> node <node>
```

```
network interface modify -vserver <vserver> -lif <fc_san_lif>
```

2. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

- Cluster quorum check
- System ID check
- Image version check
- Target platform check
- Network reachability check

The operation pauses at this stage in the network reachability check.

3. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following checks:

- Cluster health check
- Cluster LIF status check

After performing these checks, the system relocates the non-root aggregates and NAS data LIFs owned by node1 to the new controller, node3.

The controller replacement operation pauses after the resource relocation is complete.

4. Check the status of the aggregate relocation and NAS data LIF move operations:

```
system controller replace show-details
```

If the controller replacement procedure is paused, check and correct the error, if any, and then issue `resume` to continue the operation.

5. If necessary, restore and revert displaced LIFs or manually migrate and modify the node1 LIFs that failed to relocate automatically to node3.

Restore and revert displaced LIFs

- a. List any displaced LIFs:

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

- b. If any LIFs are displaced, restore the home node back to node3:

```
cluster controller-replacement network displaced-interface  
restore-home-node -node <node3_nodename> -vserver <vserver name>  
-lif-name <lif_name>
```

Manually migrate and modify LIFs

- a. Migrate the LIFs that failed to relocate automatically to node3:

```
network interface migrate -vserver <vserver name> -lif <lif_name>  
-destination-node <node3_nodename> -destination-port  
<port_on_node3>
```

- b. Modify the home node and home port for the migrated LIFs:

```
network interface modify -vserver <vserver_name> -lif  
<data_lif_name> -home-node <node3_nodename> -home-port  
<home_port>
```

6. Resume the operation to prompt the system to perform the required post-checks:

```
system controller replace resume
```

The system performs the following post-checks:

- Cluster quorum check
- Cluster health check
- Aggregates reconstruction check
- Aggregate status check
- Disk status check
- Cluster LIF status check
- Volume check

Stage 4. Relocate and retire node2

Relocate non-root aggregates and NAS data LIFs from node2 to node3

Before replacing node2 with node4, you relocate the non-root aggregates and NAS data LIFs that are owned by node2 to node3.

Before you begin

After the post-checks from the previous stage complete, the resource release for node2 starts automatically. The non-root aggregates and non-SAN data LIFs are migrated from node2 to 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.

After the aggregates and LIFs are migrated, the operation is paused for verification purposes. At this stage, you must verify whether or not all the non-root aggregates and non-SAN data LIFs are migrated to node3.



The home owner for the aggregates and LIFs are not modified; only the current owner is modified.

Steps

1. Verify that all the non-root aggregates are online and their state on node3:

```
storage aggregate show -node <node3> -state online -root false
```

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

```
cluster::> storage aggregate show -node node3 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        normal
aggr_2         825.0GB       825.0GB     0%       online   1        node2
raid_dp        normal
2 entries were displayed.
```

If the aggregates have gone offline or become foreign on node3, bring them online by using the following command on node3, once for each aggregate:

```
storage aggregate online -aggregate <aggregate_name>
```

2. Verify that all the volumes are online on node3 by using the following command on node3 and examining the output:

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

If any volumes are offline on node3, bring them online by using the following command on node3, once for each volume:

```
volume online -vserver <vserver_name> -volume <volume_name>
```

The `vserver_name` to use with this command is found in the output of the previous `volume show` command.

3. Verify that the LIFs have been moved to the correct ports and have a status of `up`. 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> -home-node <node_name> -status-admin up
```

4. If the ports currently hosting data LIFs will not exist on the new hardware, remove them from the broadcast domain:

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

5. Verify that there are no data LIFs remaining on node2 by entering the following command and examining the output:

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

Relocate failed or vetoed aggregates to node3

If any aggregates fail to relocate or are vetoed, you need to manually relocate the aggregates to node3, or if necessary, override either the vetoes or destination checks.

About this task

The system pauses the relocation operation due to the error.

Steps

1. Check the event management system (EMS) logs to determine why the aggregate failed to relocate or was vetoed.
2. Relocate any failed or vetoed aggregates:

```
storage aggregate relocation start -node <node2> -destination <node3>  
-aggregate-list <aggregate_name> -ndo-controller-upgrade true
```

3. When prompted, enter `y`.
4. You can force relocation by using one of the following methods:

Option	Description
Overriding veto checks	Use the following command: <pre>storage aggregate relocation start -node node2 -destination node3 -aggregate-list <aggregate_list> -ndo-controller-upgrade true -override-vetoes true</pre>

Option	Description
Overriding destination checks	Use the following command: <pre>storage aggregate relocation start -node node2 -destination node3 -aggregate-list <aggregate_list> -ndo-controller-upgrade true -override-vetoes true -override-destination-checks true</pre>

Retire node2

To retire node2, you shut down node2 correctly and then remove it from the rack or chassis.

Steps

1. Resume the operation:

```
system controller replace resume
```

The node halts automatically.

After you finish

After retiring node2, verify that it's powered off and not connected to the network. Leave the old hardware of node2 in this state until you complete the upgrade of both node1 and node2. You can then decommission node2 during [Stage 6. Complete the upgrade](#).

Stage 5. Install and boot node4

Install and boot node4

You install node4 in the rack, transfer node2's connections to node4, boot node4, and install ONTAP. You then reassign any of node2's spare disks, any disks belonging to the root volume, and any non-root aggregates that were not relocated to node3 earlier in the process, as outlined in this section.

About this task

The relocation operation is paused at the beginning of this stage. This process is mostly automated; the operation pauses to enable you to check its status. You must manually resume the operation.

You must netboot node4 if the ONTAP version on node4 is different to the ONTAP version on node2. After you install node4, boot it 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 by following the instructions in [Prepare for netboot](#).

Steps

1. Make sure that node4 has sufficient rack space.

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

2. Install node4 in the rack, following the instructions in the *Installation and Setup Instructions* for the node model.

3. Cable node4, moving the connections from node2 to node4.

Cable the following connections, using the instructions in the *Installation and Setup Instructions* for the node4 platform, the appropriate disk shelf document, and *HA pair management*.

Refer to [References](#) to link to *HA pair management*.

- Console (remote management port)
- Cluster and HA ports
- Data ports
- Cluster and node management ports
- Serial-attached SCSI (SAS) and Ethernet storage ports
- SAN configurations: iSCSI Ethernet, FC, and NVMe/FC switch ports



You might need to change the interconnect cables between the old and new controllers to allow interoperability between the different controller and card models. Refer to the [system installation procedures](#) for a cabling map of the Ethernet storage shelves for your systems.

For controllers introduced in ONTAP 9.15.1 and later, cluster and HA interconnects use the same ports. For switch connected configurations, it is required to connect similar ports to the same cluster switches. For example, when upgrading to an AFF A1K from an existing controller, you should connect e1a ports on both nodes to one switch and e7a ports on both nodes to the second switch.

4. 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.



When you boot node4, you might see the following warning 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.



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




Refer to [Prepare for netboot](#).

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



You must 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	<p>Configure the connection automatically by using the following command at the boot environment prompt:</p> <pre>ifconfig e0M -auto</pre>
Not running	<p>Manually configure the connection by entering the following command at the boot environment prompt:</p> <pre>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></pre> <p><i>filer_addr</i> is the IP address of the storage system (mandatory). <i>netmask</i> is the network mask of the storage system (mandatory). <i>gateway</i> is the gateway for the storage system (mandatory). <i>dns_addr</i> is the IP address of a name server on your network (optional). <i>dns_domain</i> is the DNS domain name (optional).</p> <div>  <p>Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p> </div>

7. Perform netboot on node4:

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

The `<path_to_the_web-accessible_directory>` should lead to where you downloaded the `<ontap_version>_image.tgz` in Step 1 in the section [Prepare for netboot](#).



Do not interrupt the boot.

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

This menu option downloads and installs the new 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 ONTAP, and not upgrades of controllers.



Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the incorrect image might install. This issue applies to all ONTAP releases. The netboot procedure combined with option (7) `Install new software` wipes the boot media and places the same ONTAP version on both image partitions.

9. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL:

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

10. Complete the following substeps to reboot the controller module:

- 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}
```

- 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 must be restored.

11. Select maintenance mode 5 from the boot menu and enter `y` when you are prompted to continue with the boot.
12. Verify that the controller and chassis are configured as HA:

```
ha-config show
```

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

```
Chassis HA configuration: ha  
Controller HA configuration: ha
```



System records 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.

13. If the controller and chassis are not configured as HA, use the following commands to correct the configuration:

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

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


14. Confirm that all Ethernet ports used to connect to the Ethernet shelves are configured as storage:

```
storage port show
```

The output displayed depends on the system configuration. The following output example is for a node with a single storage card in slot11. The output for your system might be different:

```
*> storage port show
Port Type Mode      Speed (Gb/s) State      Status  VLAN ID
---- -
e11a ENET storage 100 Gb/s    enabled  online  30
e11b ENET storage 100 Gb/s    enabled  online  30
```

15. Modify the ports that are not set to storage:

```
storage port modify -p <port> -m storage
```

All Ethernet ports connected to storage shelves must be configured as storage to allow access to the disks and shelves.

16. Exit maintenance mode:

```
halt
```

Interrupt the autoboot by pressing Ctrl-C at the boot environment prompt.

17. On node3, check the system date, time, and time zone:

```
date
```

18. On node4, check the date by using the following command at the boot environment prompt:

```
show date
```

19. If necessary, set the date on node4:

```
set date <mm/dd/yyyy>
```

20. On node4, check the time by using the following command at the boot environment prompt:

```
show time
```

21. If necessary, set the time on node4:

```
set time <hh:mm:ss>
```

22. In boot loader, set the partner system ID on node4:

```
setenv partner-sysid <node3_sysid>
```

For node4, partner-sysid must be that of node3.

Save the settings:

```
saveenv
```

23. Verify the `partner-sysid` for node4:

```
printenv partner-sysid
```

24. If you have NetApp Storage Encryption (NSE) drives installed, perform the following steps.



If you have not already done so earlier in the procedure, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use.

- a. Set `bootarg.storageencryption.support` to `true` or `false`.

If the following drives are in use...	Then...
NSE drives that conform to FIPS 140-2 Level 2 self-encryption requirements	<code>setenv bootarg.storageencryption.support true</code>
NetApp non-FIPS SEDs	<code>setenv bootarg.storageencryption.support false</code>

- b. Go to the special boot menu and select option (10) `Set Onboard Key Manager recovery secrets`.

Enter the passphrase and the backup information that you recorded earlier procedure. See [Manage storage encryption using the Onboard Key Manager](#).

25. Boot the node into boot menu:

```
boot_ontap menu
```

26. When node4 stops at the boot menu, reassign the node2 disks to node4 by running the following command on node4:

```
boot_after_controller_replacement
```

After a short delay, you are prompted to enter the name of the node that is being replaced. If there are shared disks (also called Advanced Disk Partitioning (ADP) or partitioned disks), you are prompted to enter the node name of the HA partner.

These prompts might get buried in the console messages. If you do not enter a node name or enter an incorrect name, you are prompted to enter the name again.

Expand the console output example

```
LOADER-A> boot_ontap menu
.
.
<output truncated>
.
All rights reserved.
*****
*                                     *
* Press Ctrl-C for Boot Menu. *
*                                     *
*****
.
<output truncated>
.
Please choose one of the following:
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 22/7
(22/7)                                     Print this secret List
(25/6)                                     Force boot with multiple filesystem
disks missing.
(25/7)                                     Boot w/ disk labels forced to clean.
(29/7)                                     Bypass media errors.
(44/4a)                                    Zero disks if needed and create new
flexible root volume.
(44/7)                                     Assign all disks, Initialize all
disks as SPARE, write DDR labels
.
.
<output truncated>
.
.
(wipeconfig)                             Clean all configuration on boot
device
(boot_after_controller_replacement) Boot after controller upgrade
```

```

(boot_after_mcc_transition)      Boot after MCC transition
(9a)                             Unpartition all disks and remove
their ownership information.
(9b)                             Clean configuration and
initialize node with partitioned disks.
(9c)                             Clean configuration and
initialize node with whole disks.
(9d)                             Reboot the node.
(9e)                             Return to main boot menu.
The boot device has changed. System configuration information could
be lost. Use option (6) to
restore the system configuration, or option (4) to initialize all
disks and setup a new system.
Normal Boot is prohibited.
Please choose one of the following:
(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? boot_after_controller_replacement
This will replace all flash-based configuration with the last backup
to disks. Are you sure
you want to continue?: yes
.
.
<output truncated>
.
.
Controller Replacement: Provide name of the node you would like to
replace:
<nodename of the node being replaced>
Controller Replacement: Provide High Availability partner of node1:
<nodename of the partner of the node being replaced>
Changing sysid of node node2 disks.
Fetched sanown old_owner_sysid = 536940063 and calculated old sys id
= 536940063
Partner sysid = 4294967295, owner sysid = 536940063
.
.

```

```

<output truncated>
.
.
varfs_backup_restore: restore using /mroot/etc/varfs.tgz
varfs_backup_restore: attempting to restore /var/kmip to the boot
device
varfs_backup_restore: failed to restore /var/kmip to the boot device
varfs_backup_restore: attempting to restore env file to the boot
device
varfs_backup_restore: successfully restored env file to the boot
device wrote
    key file "/tmp/rndc.key"
varfs_backup_restore: timeout waiting for login
varfs_backup_restore: Rebooting to load the new varfs
Terminated
<node reboots>
System rebooting...
.
.
Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...
Rebooting to load the restored env file...
.
System rebooting...
.
.
.
<output truncated>
.
.
.
.
WARNING: System ID mismatch. This usually occurs when replacing a
boot device or NVRAM cards!
Override system ID? {y|n} y
.
.
.
.
Login:

```



In the above console output example, ONTAP will prompt you for the partner node name if the system uses Advanced Disk Partitioning (ADP) disks.

27. At the LOADER prompt, boot up:

```
boot_ontap menu
```

Now, on booting, the node can detect all the disks that were previously assigned to it and can boot up as expected.

When the cluster nodes you are replacing use root volume encryption, ONTAP is unable to read the volume information from the disks. Restore the keys for the root volume:

If the root volume is encrypted, recover the onboard key-management secrets so the system can find the root volume.

a. Return to the special boot menu:

```
LOADER> boot_ontap menu
```

```
Please choose one of the following:
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.

Selection (1-11)? 10
```

b. Select **(10) Set Onboard Key Manager recovery secrets**

c. Enter **y** at the following prompt:

```
This option must be used only in disaster recovery procedures. Are you sure?
(y or n): y
```

d. At the prompt, enter the key-manager passphrase.

e. Enter the backup data when prompted.



You must have obtained the passphrase and backup data in the [Prepare the nodes for upgrade](#) section of this procedure.

f. After the system boots to the special boot menu again, run option **(1) Normal Boot**



You might encounter an error at this stage. If an error occurs, repeat the substeps in [Step 27](#) until the system boots normally.

Verify the node4 installation

You must verify that the physical ports from node2 map correctly to the physical ports on node4. This will enable node4 to communicate with other nodes in the cluster and with the network after the upgrade.

About this task

Refer to [References](#) to link to the *Hardware Universe* to capture information about the ports on the new nodes. You will use the information later in this section.

Physical port layout might vary, depending on the model of the nodes. When the new node boots up, ONTAP will try to determine which ports should host cluster LIFs in order to automatically come into quorum.

If the physical ports on node2 do not map directly to the physical ports on node4, the subsequent section [Restore network configuration on node4](#) must be used to repair network connectivity.

After you install and boot node4, you must verify that it is installed correctly. You must wait for node4 to join quorum and then resume the relocation operation.

At this point in the procedure, the operation will have paused as node4 joins quorum.

Steps

1. Verify that node4 has joined quorum:

```
cluster show -node node4 -fields health
```

The output of the `health` field should be `true`.

2. Verify that node4 is part of the same cluster as node3 and that it is healthy:

```
cluster show
```

3. Switch to advanced privilege mode:

```
set advanced
```

4. Check the status of the controller replacement operation and verify that it is in a paused state and in the same state it was in before node2 was halted to perform the physical tasks of installing new controllers and moving cables:

```
system controller replace show
```

```
system controller replace show-details
```

5. Resume the controller replacement operation:

```
system controller replace resume
```

6. Controller replacement will pause for intervention with the following message:

```

Cluster::*> system controller replace show
Node                Status                Error-Action
-----
Node2(now node4) Paused-for-intervention  Follow the instructions
given in
Step Details
Node2
Step Details:
-----
To complete the Network Reachability task, the ONTAP network
configuration must be
manually adjusted to match the new physical network configuration of the
hardware.
This includes:

1. Re-create the interface group, if needed, before restoring VLANs. For
detailed
commands and instructions, refer to the "Re-creating VLANs, ifgrps, and
broadcast
domains" section of the upgrade controller hardware guide for the ONTAP
version
running on the new controllers.
2. Run the command "cluster controller-replacement network displaced-
vlans show"
to check if any VLAN is displaced.
3. If any VLAN is displaced, run the command "cluster controller-
replacement
network displaced-vlans restore" to restore the VLAN on the desired
port.
2 entries were displayed.

```



In this procedure, section *Re-creating VLANs, ifgrps, and broadcast domains* has been renamed *Restoring network configuration on node4*.

7. With the controller replacement in a paused state, proceed to the next section of this document to restore network configuration on the node.

Restore network configuration on node4

After you confirm that node4 is in quorum and can communicate with node3, verify that node2's VLANs, interface groups and broadcast domains are seen on node4. Also, verify that all node4 network ports are configured in their correct broadcast domains.

About this task

For more information on creating and re-creating VLANs, interface groups, and broadcast domains, refer to [References](#) to link to *Network Management*.

Steps

1. List all the physical ports that are on upgraded node2 (referred to as node4):

```
network port show -node node4
```

All physical network ports, VLAN ports and interface group ports on the node are displayed. From this output you can see any physical ports that have been moved into the `Cluster` broadcast domain by ONTAP. You can use this output to aid in deciding which ports should be used as interface group member ports, VLAN base ports or standalone physical ports for hosting LIFs.

2. List the broadcast domains on the cluster:

```
network port broadcast-domain show
```

3. List the network port reachability of all ports on node4:

```
network port reachability show
```

The output from the command looks similar to the following example:

```
ClusterA::*> network port reachability show
```

Node	Port	Expected Reachability	Reachability

node1_node3			
	e0M	Default:Mgmt	ok
	e10a	Default:Default	ok
	e10b	-	no-reachability
	e10c	Default:Default	ok
	e10d	-	no-reachability
	e1a	Cluster:Cluster	ok
	e1b	-	no-reachability
	e7a	Cluster:Cluster	ok
	e7b	-	no-reachability
node2_node4			
	e0M	Default:Mgmt	ok
	e10a	Default:Default	ok
	e10b	-	no-reachability
	e10c	Default:Default	ok
	e10d	-	no-reachability
	e1a	Cluster:Cluster	ok
	e1b	-	no-reachability
	e7a	Cluster:Cluster	ok
	e7b	-	no-reachability

18 entries were displayed.

In the above example, node2_node4 is just booted after controller replacement. It has several ports that have no reachability and are pending a reachability scan.

4. Repair the reachability for each of the ports on node4 with a reachability status other than `ok`. Run the following command, first on any physical ports, then on any VLAN ports, one at a time:

```
network port reachability repair -node <node_name> -port <port_name>
```

The output looks like the following example:

```
Cluster ::> reachability repair -node node2_node4 -port e10a
```

```
Warning: Repairing port "node2_node4: e10a" may cause it to move into a
different broadcast domain, which can cause LIFs to be re-homed away
from the port. Are you sure you want to continue? {y|n}:
```

A warning message, as shown above, is expected for ports with a reachability status that might be different

from the reachability status of the broadcast domain where it is currently located.

Review the connectivity of the port and answer `y` or `n` as appropriate.

Verify that all physical ports have their expected reachability:

```
network port reachability show
```

As the reachability repair is performed, ONTAP attempts to place the ports in the correct broadcast domains. However, if a port's reachability cannot be determined and does not belong to any of the existing broadcast domains, ONTAP will create new broadcast domains for these ports.

5. If interface group configuration does not match the new controller physical port layout, modify it by using the following steps.
 - a. You must first remove physical ports that should be interface group member ports from their broadcast domain membership. You can do this by using the following command:

```
network port broadcast-domain remove-ports -broadcast-domain  
<broadcast_domain_name> -ports <node_name:port_name>
```

- b. Add a member port to an interface group:

```
network port ifgrp add-port -node <node_name> -ifgrp <ifgrp> -port  
<port_name>
```

- c. The interface group is automatically added to the broadcast domain about a minute after the first member port is added.
 - d. Verify that the interface group was added to the appropriate broadcast domain:

```
network port reachability show -node <node_name> -port <ifgrp>
```

If the interface group's reachability status is not `ok`, assign it to the appropriate broadcast domain:

```
network port broadcast-domain add-ports -broadcast-domain  
<broadcast_domain_name> -ports <node:port>
```

6. Assign appropriate physical ports to the `Cluster` broadcast domain:

- a. Determine which ports have reachability to the `Cluster` broadcast domain:

```
network port reachability show -reachable-broadcast-domains Cluster:Cluster
```

- b. Repair any port with reachability to the `Cluster` broadcast domain, if its reachability status is not `ok`:

```
network port reachability repair -node <node_name> -port <port_name>
```

7. Move the remaining physical ports into their correct broadcast domains by using one of the following commands:

```
network port reachability repair -node <node_name> -port <port_name>
```

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

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

Verify that there are no unreachable or unexpected ports present. Check the reachability status for all physical ports by using the following command and examining the output to confirm the status is ok:

```
network port reachability show -detail
```

8. Restore any VLANs that might have become displaced by using the following steps:

a. List displaced VLANs:

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

Output like the following should display:

```
Cluster::*> displaced-vlans show
(cluster controller-replacement network displaced-vlans show)

      Original
Node      Base Port      VLANs
-----
Node1     a0a            822, 823
          e10a            822, 823
```

b. Restore VLANs that were displaced from their previous base ports:

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

The following is an example of restoring VLANs that have been displaced from interface group a0a back onto the same interface group:

```
Cluster::*> displaced-vlans restore -node node2_node4 -port a0a
-destination-port a0a
```

The following is an example of restoring displaced VLANs on port "e10a" to "e10b":

```
Cluster::*> displaced-vlans restore -node node2_node4 -port e10a
-destination-port e10b
```

When a VLAN restore is successful, the displaced VLANs are created on the specified destination port. The VLAN restore fails if the destination port is a member of an interface group, or if the destination port is down.

Wait about one minute for newly restored VLANs to be placed into their appropriate broadcast domains.

c. Create new VLAN ports as needed for VLAN ports that are not in the `cluster controller-replacement network displaced-vlans show` output but should be configured on other physical ports.

9. Delete any empty broadcast domains after all port repairs have been completed:

```
network port broadcast-domain delete -broadcast-domain <broadcast_domain_name>
```

10. Verify port reachability:

```
network port reachability show
```

When all ports are correctly configured and added to the correct broadcast domains, the `network port reachability show` command should report the reachability status as `ok` for all connected ports, and the status as `no-reachability` for ports with no physical connectivity. If any ports report a status other than these two, perform the reachability repair and add or remove ports from their broadcast domains as instructed in [Step 4](#).

11. Verify that all ports have been placed into broadcast domains:

```
network port show
```

12. Verify that all ports in the broadcast domains have the correct maximum transmission unit (MTU) configured:

```
network port broadcast-domain show
```

13. Restore LIF home ports, specifying the Vserver(s) and LIF(s) home ports, if any, that need to be restored:

a. List any LIFs that are displaced:

```
displaced-interface show
```

b. Restore LIF home ports:

```
displaced-interface restore-home-node -node <node_name> -vserver  
<vserver_name> -lif-name <LIF_name>
```

14. Verify that all LIFs have a home port and are administratively up:

```
network interface show -fields home-port, status-admin
```

Restore key-manager configuration on node4

If you are using NetApp Volume Encryption (NVE) and NetApp Aggregate Encryption (NAE) to encrypt volumes on the system you are upgrading, the encryption configuration must be synchronized to the new nodes. If you do not synchronize the key-manager, when you relocate the node2 aggregates from node3 to node4 by using ARL, failures might occur because node4 does not have the required encryption keys to bring encrypted volumes and aggregates online.

About this task

Synchronize the encryption configuration to the new nodes by performing the following steps:

Steps

1. Run the following command from node4:

```
security key-manager onboard sync
```

2. Verify that the SVM-KEK key is restored to "true" on node4 before you relocate the data aggregates:

```
::> security key-manager key query -node node4 -fields restored -key  
-type SVM-KEK
```

Example

```
::> security key-manager key query -node node4 -fields restored -key  
-type SVM-KEK
```

node	vserver	key-server	key-id
restored			
-----	-----	-----	-----
node4	svm1	""	000000000000000000200000000000a008a81976
true			2190178f9350e071fbb90f00000000000000000

Move non-root aggregates and NAS data LIFs owned by node2 from node3 to node4

After you verify network configuration on node4 and before you relocate aggregates from node3 to node4, you must verify that the NAS data LIFs belonging to node2 that are currently on node3 are relocated from node3 to node4. You must also verify that the SAN LIFs exist 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 will verify that the LIFs are healthy and located on appropriate ports after you bring node4 online.

Steps

1. The iSCSI LIFs automatically find the correct home ports using the reachability scan. The FC and NVMe/FC SAN LIFs do not move automatically. They continue to show the home port they were on before upgrading.

Check the SAN LIFs on node4:

- a. Modify any iSCSI SAN LIFs reporting a "down" operation status to the new data ports:

```
network interface modify -vserver <vserver> -lif <iscsi_san_lif>  
admin down
```

```
network interface modify -vserver <vserver> -lif <iscsi_san_lif> port  
<new_port> node <node>
```

```
network interface modify -vserver <vserver> -lif <iscsi_san_lif>
```

- b. Modify any FC and NVMe/FC SAN LIFs that are home to the new controller and reporting a "down" operational status to the FCP ports on the new controller:

```
network interface modify -vserver <vserver> -lif <fc_san_lif> admin  
down
```

```
network interface modify -vserver <vserver> -lif <fc_san_lif> port  
<new_port> node <node>
```

```
network interface modify -vserver <vserver> -lif <fc_san_lif>
```

2. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

- Cluster quorum check
- System ID check
- Image version check
- Target platform check
- Network reachability check

The operation pauses at this stage in the network reachability check.

3. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following checks:

- Cluster health check
- Cluster LIF status check

After performing these checks, the system relocates the non-root aggregates and NAS data LIFs owned by

node2 to the new controller, node4.

The controller replacement operation pauses after the resource relocation is complete.

4. Check the status of the aggregate relocation and NAS data LIF move operations:

```
system controller replace show-details
```

If the controller replacement procedure is paused, check and correct the error, if any, and then issue `resume` to continue the operation.

5. If necessary, restore and revert displaced LIFs or manually migrate and modify the node2 LIFs that failed to relocate automatically to node4.

Restore and revert displaced LIFs

- a. List any displaced LIFs:

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

- b. If any LIFs are displaced, restore the home node back to node4:

```
cluster controller-replacement network displaced-interface  
restore-home-node -node <node4_nodename> -vserver <vserver name>  
-lif-name <lif_name>
```

Manually migrate and modify LIFs

- a. Migrate the LIFs that failed to relocate automatically to node4:

```
network interface migrate -vserver <vserver name> -lif <lif_name>  
-destination-node <node4_nodename> -destination-port  
<port_on_node4>
```

- b. Modify the home node and home port for the migrated LIFs:

```
network interface modify -vserver <vserver_name> -lif  
<data_lif_name> -home-node <node4_nodename> -home-port  
<home_port>
```

6. Resume the operation to prompt the system to perform the required post-checks:

```
system controller replace resume
```


The system performs the following post-checks:

- Cluster quorum check
- Cluster health check
- Aggregates reconstruction check
- Aggregate status check
- Disk status check
- Cluster LIF status check
- Volume check

Stage 6. Complete the upgrade

Manage authentication using KMIP servers

You can use Key Management Interoperability Protocol (KMIP) servers to manage authentication keys.

Steps

1. Add a new controller:

```
security key-manager external enable
```

2. Add the key manager:

```
security key-manager external add-servers -key-servers  
key_management_server_ip_address
```

3. Verify that the key management servers are configured and available to all nodes in the cluster:

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

4. Restore the authentication keys from all linked key management servers to the new node:

```
security key-manager external restore -node new_controller_name
```

Confirm that the new controllers are set up correctly

To confirm 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 confirm that node3 owns node1's aggregates and that node4 owns node2's aggregates, and that the volumes for both nodes are online.

Steps

1. After the post-checks of node2, the storage failover and cluster HA pair for the node2 cluster are enabled. When the operation is done, both nodes show as completed and the system performs some cleanup operations.
2. Verify that storage failover is enabled:

```
storage failover show
```

The following example shows the output of the command when storage failover is enabled:

```
cluster::> storage failover show
```

		Takeover	
Node	Partner	Possible	State Description
node3	node4	true	Connected to node4
node4	node3	true	Connected to node3

3. Verify that node3 and node4 belong to the same cluster by using the following command and examining the output:

```
cluster show
```

4. Verify that node3 and node4 can access each other's storage by using the following command and examining the output:

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

5. Verify that neither node3 nor node4 owns data LIFs home-owned by other nodes in the cluster by using the following command and examining the output:

```
network interface show
```

If neither node3 or node4 owns data LIFs home-owned by other nodes in the cluster, revert the data LIFs to their home owner:

```
network interface revert
```

6. Verify that node3 owns the aggregates from node1 and that node4 owns the aggregates from node2:

```
storage aggregate show -owner-name <node3>
```

```
storage aggregate show -owner-name <node4>
```

7. Determine whether any volumes are offline:

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

```
volume show -node <node4> -state offline
```

8. If any volumes are offline, compare them with the list of offline volumes that you captured in the section [Prepare the nodes for upgrade](#), and bring online any of the offline volumes, as required, by using the following command, once for each volume:

```
volume online -vserver <vserver_name> -volume <volume_name>
```

9. Install new licenses for the new nodes by using the following command 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 multiple licenses at once, separating each license key by a comma.

10. Remove all of the old licenses from the original nodes by using one of the following commands:

```
system license clean-up -unused -expired
```

```
system license delete -serial-number <node_serial_number> -package  
<licensable_package>
```

- Delete all expired licenses:

```
system license clean-up -expired
```

- Delete all unused licenses:

```
system license clean-up -unused
```

- Delete a specific license from a cluster by using 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.

11. Verify that the licenses are correctly installed by using the following command and examining the output:

```
system license show
```

You can compare the output with the output that you captured in the section [Prepare the nodes for upgrade](#).

12. If self-encrypting drives are being used in the configuration and you have set the `kmip.init.maxwait` variable to `off` (for example, in [Install and boot node4, Step 24](#)), you must unset the variable:

```
set diag; systemshell -node <node_name> -command sudo kenv -u -p  
kmip.init.maxwait
```

13. Configure the SPs by using the following command on both nodes:

```
system service-processor network modify -node <node_name>
```

Refer to [References](#) to link to the *System Administration Reference* for information about the SPs and the *ONTAP 9.8 Commands: Manual Page Reference* for detailed information about the `system service-processor network modify` command.

14. If you want to set up a switchless cluster on the new nodes, refer to [References](#) to link to the *NetApp Support Site* and follow the instructions in *Transitioning to a two-node switchless cluster*.

After you finish

If Storage Encryption is enabled on node3 and node4, complete the section [Set up Storage Encryption on the new controller module](#). Otherwise, complete the section [Decommission the old system](#).

Set 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. You must 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 external show-status
```

```
security key-manager onboard show-backup
```

2. Add the key management servers listed in the previous step to the key management server list in the new controller.
 - a. Add the key management server:

```
security key-manager external add-servers -key-servers  
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 external 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 external enable
```

- 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 external restore -node new_controller_name
```

Set up NetApp Volume or Aggregate Encryption on the new controller module

If the replaced controller or high availability (HA) partner of the new controller uses NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE), you must configure the new controller module for NVE or NAE.

About this task

This procedure includes steps that are performed on the new controller module. You must enter the command on the correct node.

Onboard Key Manager

Configure NVE or NAE using the Onboard Key Manager.

Steps

1. Restore authentication keys from all linked key management servers to the new node:

```
security key-manager onboard sync
```

External Key Management

Configure NVE or NAE using External Key Management.

Steps

1. Verify that the key management servers are still available, their status, and their authentication key information:

```
security key-manager key query -node node
```

2. Add the key management servers listed in the previous step to the key management server list in the new controller:

- a. Add the key management server:

```
security key-manager external add-servers -key-servers  
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 external 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 external enable
```

- 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 external restore
```

This command needs the OKM passphrase

For more information, see the Knowledge Base article [How to restore external key manager server configuration from the ONTAP boot menu](#).

After you finish

Check if any volumes were taken offline because authentication keys were not available or EKM servers could not be reached. Bring those volumes back online by using the `volume online` command.

Decommission 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. Refer to [References](#) to link to the *NetApp Support Site* and log in.
2. Select **Products > My Products** from the menu.
3. On the **View Installed Systems** page, choose which **Selection Criteria** you want to use to display information about your system.

You can choose one of the following to locate your system:

- Serial Number (located on the back of the unit)
- Serial Numbers for My Location

4. Select **Go!**

A table displays cluster information, including the serial numbers.

5. Locate the cluster in the table and select **Decommission this system** from the Product Tool Set drop-down menu.

Resume SnapMirror operations

You can resume SnapMirror transfers that were quiesced before upgrade and resume the SnapMirror relationships. The updates are on schedule after the upgrade is completed.

Steps

1. Verify the SnapMirror status on the destination:

```
snapmirror show
```

2. Resume the SnapMirror relationship:

```
snapmirror resume -destination-vserver vservers_name
```

Troubleshoot

Aggregate relocation failures

Aggregate relocation (ARL) might fail at different points during the upgrade.

Check for aggregate relocation failure

During the procedure, ARL might fail in Stage 2, Stage 3, or Stage 5.

Steps

1. Enter the following command and examine the output:

```
storage aggregate relocation show
```

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, refer to [References](#) to link to the *ONTAP 9.8 Commands: Manual Page Reference*.

Aggregates originally on node1 are owned by node4 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 correctly, 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 [Prepare the nodes for 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 [Prepare the nodes for upgrade](#) and note any aggregates that were not correctly relocated.
3. Relocate the aggregates left behind on node4:

```
storage aggregate relocation start -node node4 -aggr aggr_node_1 -destination node3
```

Do not use the `-ndo-controller-upgrade` parameter during this relocation.

4. 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 the pre-check phase

Node1 or node2 crashes before the pre-check phase with HA pair still enabled

If either node1 or node2 crashes before the pre-check phase, no aggregates have been relocated yet and the HA pair 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.

Reboots, panics, or power cycles during first resource-release phase

Node1 crashes during the first resource-release phase with HA pair still enabled

Some or all aggregates have been relocated from node1 to node2, and HA pair is still enabled. Node2 takes over node1's root volume and any non-root aggregates that were not relocated.

About this task

Ownership of aggregates that were relocated look the same as the ownership of non-root aggregates that were taken over because the home owner has not changed.

When node1 enters the `waiting for giveback` state, node2 gives back all of the node1 non- root

aggregates.

Steps

1. After node1 is booted up, all the non-root aggregates of node1 have moved back to node1. You must perform a manual aggregate relocation of the aggregates from node1 to node2:
`storage aggregate relocation start -node node1 -destination node2 -aggregate -list * -ndocontroller-upgrade true`
2. Continue with the node-pair upgrade procedure.

Node1 crashes during the first resource-release phase while HA pair is disabled

Node2 does 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.

Node2 fails during the first resource-release phase with HA pair still enabled

Node1 has relocated some or all of its aggregates to node2. The HA pair is enabled.

About this task

Node1 takes over all of node2's aggregates as well as any of its own aggregates that it had relocated to node2. When node2 boots up, the aggregate relocation is completed automatically.

Steps

1. Bring up node2.
2. Continue with the node-pair upgrade procedure.

Node2 crashes during the first resource-release phase and after HA pair is disabled

Node1 does not take over.

Steps

1. Bring up node2.

A client outage occurs for all aggregates while node2 is booting up.
2. Continue with the rest of the node-pair upgrade procedure.

Reboots, panics, or power cycles during the first verification phase

Node2 crashes during the first verification phase with HA pair disabled

Node3 does not take over following a node2 crash as the HA pair is already disabled.

Steps

1. Bring up node2.

A client outage occurs for all aggregates while node2 is booting up.
2. Continue with the node-pair upgrade procedure.

Node3 crashes during the first verification phase with HA pair disabled

Node2 does 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.

Reboots, panics, or power cycles during first resource-regain phase

Node2 crashes during the first resource-regain phase during aggregate relocation

Node2 has relocated some or all of its aggregates from node1 to node3. Node3 serves data from aggregates that were relocated. The HA pair is disabled and hence there is no takeover.

About this task

There is client outage for aggregates that were not relocated. On booting up node2, the aggregates of node1 are relocated to node3.

Steps

1. Bring up node2.
2. Continue with the node-pair upgrade procedure.

Node3 crashes during the first resource-regain phase during aggregate relocation

If node3 crashes while node2 is relocating aggregates to node3, the task continues after node3 boots up.

About this task

Node2 continues to serve remaining aggregates, but aggregates that were already relocated to node3 encounter client outage while node3 is booting up.

Steps

1. Bring up node3.
2. Continue with the controller upgrade.

Reboots, panics, or power cycles during post-check phase

Node2 or node3 crashes during the post-check phase

The HA pair is disabled hence this is no takeover. There is a client outage for aggregates belonging to the node that rebooted.

Steps

1. Bring up the node.
2. Continue with the node-pair upgrade procedure.

Reboots, panics, or power cycles during second resource-release phase

Node3 crashes during the second resource-release phase

If node3 crashes while node2 is relocating aggregates, the task continues after node3 boots up.

About this task

Node2 continues to serve remaining aggregates but aggregates that were already relocated to node3 and node3's own aggregates encounter client outages while node3 is booting.

Steps

1. Bring up node3.
2. Continue with the controller upgrade procedure.

Node2 crashes during the second resource-release phase

If node2 crashes during aggregate relocation, node2 is not taken over.

About this task

Node3 continues to serve the aggregates that have been relocated, but the aggregates owned by node2 encounter client outages.

Steps

1. Bring up node2.
2. Continue with the controller upgrade procedure.

Reboots, panics, or power cycles during the second verification phase**Node3 crashes during the second verification phase**

If node3 crashes during this phase, takeover does not happen because the HA pair is already disabled.

About this task

There is a client outage for all aggregates until node3 reboots.

Steps

1. Bring up node3.
2. Continue with the node-pair upgrade procedure.

Node4 crashes during the second verification phase

If node4 crashes during this phase, takeover does not happen. Node3 serves data from the aggregates.

About this task

There is an outage for non-root aggregates that were already relocated until node4 reboots.

Steps

1. Bring up node4.
2. Continue with the node-pair upgrade procedure.

Issues that can arise in multiple stages of the procedure

Some issues can occur during 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	
		Possible	State Description
node1	node2	false	Unknown
node2	node1	false	Node owns partner aggregates as part of the non-disruptive head upgrade procedure. Takeover is not possible: Storage failover is disabled.

The following example shows the output of the `storage failover show` command after a reboot or panic:

```
cluster::> storage failover show
```

Node	Partner	Takeover	
		Possible	State Description
node1	node2	-	Unknown
node2	node1	false	Waiting for node1, Partial giveback, Takeover is not possible: Storage failover is disabled

Although the output says that a node is in partial giveback and that storage failover is disabled, you can disregard this message.

Steps

No action is required; continue with the node-pair upgrade procedure.

LIF migration failure

After you migrate LIFs, they might not come online after migration in Stage 2, Stage 3, or Stage 5.

Steps

1. Verify that the 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.

2. Check the physical connectivity of the network cable if the physical state of the port is down.

References

When performing the procedures in this content, you might need to consult reference content or go to reference websites.

Reference content

Content specific to this upgrade are listed in the table below.

Content	Description
Administration overview with the CLI	Describes how to administer ONTAP systems, shows you how to use the CLI interface, how to access the cluster, how to manage nodes, and much more.
Decide whether to use System Manager or the ONTAP CLI for cluster setup	Describes how to set up and configure ONTAP.
Disk and aggregate management with the CLI	Describes how to manage ONTAP physical storage using the CLI. It shows you how to create, expand, and manage aggregates, how to work with Flash Pool aggregates, how to manage disks, and how to manage RAID policies.
HA pair management	Describes how to install and manage high-availability clustered configurations, including storage failover and takeover/giveback.
Logical storage management with the CLI	Describes how to efficiently manage your logical storage resources, using volumes, FlexClone volumes, files, and LUNs, FlexCache volumes, deduplication, compression, qtrees, and quotas.
MetroCluster Upgrade and Expansion	Provides procedures for upgrading controller and storage models in the MetroCluster configuration, transitioning from a MetroCluster FC to a MetroCluster IP configuration, and expanding the MetroCluster configuration by adding additional nodes.
Network Management	Describes how to configure and manage physical and virtual network ports (VLANs and interface groups), LIFs, routing, and host-resolution services in clusters; optimize network traffic by load balancing; and monitor the cluster by using SNMP.
ONTAP 9.13.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.13.1 commands.
ONTAP 9.14.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.14.1 commands.
ONTAP 9.15.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.15.1 commands.

Content	Description
SAN management with the CLI	Describes how to configure and manage LUNs, igroups, and targets using the iSCSI and FC protocols, and namespaces and subsystems using the NVMe/FC protocol.
SAN configuration reference	Contains information about FC and iSCSI topologies and wiring schemes.
Upgrade by moving volumes or storage	Describes how to quickly upgrade controller hardware in a cluster by moving storage or volumes. Also describes how to convert a supported model to a disk shelf.
Upgrade ONTAP	Contains instructions for downloading and upgrading ONTAP.
Use "system controller replace" commands to upgrade controller models in the same chassis	Describes the aggregate relocation procedures needed to non-disruptively upgrade a system, keeping the old system chassis and disks.
Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers running ONTAP 9.8 by using "system controller replace" commands.
Use aggregate relocation to manually upgrade controller hardware running ONTAP 9.8 or later	Describes the aggregate relocation procedures needed to perform manual non-disruptive controller upgrades running ONTAP 9.8 or later.

Reference sites

The [NetApp Support Site](#) also contains documentation about network interface cards (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.

Access [ONTAP 9 documentation](#).

Access the [Active IQ Config Advisor](#) tool.

Use "system controller replace" commands to upgrade controller models in the same chassis

Learn about this ARL upgrade procedure

There are several aggregate relocation (ARL) methods for upgrading controller hardware. This procedure describes how to upgrade storage controllers in an HA pair by converting the existing system to the replacement system, keeping the existing system chassis and disks.

ARL takes advantage of the HA configuration and cluster interconnect communication. This allows 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 the non-root aggregates between the old controller nodes. After installing the replacement nodes, you migrate the non-root aggregates from the old controller nodes to the replacement controller nodes. The data hosted on the nodes that you are upgrading is

accessible during the upgrade procedure. You also migrate data LIFs between nodes in the cluster as you proceed.

The controller hardware that you replace depends on the existing system model type:

If your existing system is...	Then...
AFF A250, AFF C250	Swap the two AFF A250 or AFF C250 controllers with the new controllers and I/O modules.
AFF A800, AFF C800	Swap the two AFF A800 or AFF C800 controllers with the new controllers and I/O modules.
AFF A220, AFF A200, AFF C190, FAS2620, or FAS2720	Swap the controller module on each node on the old controller with the new module.
AFF A700 or FAS9000	Swap the controller and NVRAM modules on each node on the old controller with the new modules. Note: You don't need to move, disconnect, or reconnect the I/O cards, data cables, disk shelves, and disks.



The terms **node1** and **node2**, are used only as a reference to node names in this document. When following the procedure, you must substitute the actual names of your nodes.

Requirements and limitations

You need to consider important factors before you start your upgrade procedure.



You must review all the following important information before starting the upgrade procedure.

- This procedure is complex and assumes that you have advanced ONTAP administration skills. You also should read and understand the [guidelines for upgrading controllers](#) and [ARL upgrade sequence](#) before beginning the upgrade.
- This procedure assumes that the replacement controller hardware is new and has not been used in another system. 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 as part of another ONTAP cluster or as a standalone single node system.
- 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.
- If you have a switch that is not supported by the ONTAP version and the replacement system that you are upgrading to, refer to [References](#) to link to the *Hardware Universe*.
- The AFF A250 and AFF C250 systems use 10/25 GbE onboard ports for cluster interconnect. To upgrade two-node switchless clusters of an AFF A250 or AFF C250 system to an AFF A50, AFF A30, AFF C60, or AFF C30 system, you must use an X60132A, 4-port 10/25 GbE card in slot1 on both nodes. This provides cluster interconnect for cluster LIFs to a temporary port on the target platform.
- The AFF A30, AFF A50, AFF A70, AFF A90, AFF A1K, AFF C30, AFF C60, AFF C80, FAS70, and FAS90 systems share 100GbE network ports for both cluster and HA connections. These systems can support 10GbE or 25GbE cluster connections to legacy cluster switches; however, NetApp recommends updating to 100GbE cluster speeds when the 10GbE and 25GbE switches are no longer required. For more information, see the following Knowledge Base articles:

- [How to configure 10G or 25G cluster ports on a new cluster setup](#)
- [How to convert existing 10G or 25G cluster ports to 40G or 100G cluster ports](#)

If you cannot link up e0a or e0b cluster ports on the existing node to the cluster ports on the new node, see the following for more information:

- [NetApp Bugs Online Bug ID CONTAP-166978](#)
- [NetApp Bugs Online Bug ID 1127315](#)
- The ASA A900, AFF A900, and FAS9500 systems only support high-line power (200V to 240V). If your AFF A700 or FAS9000 system is running on low-line power (100V to 120V), you must convert the AFF A700 or FAS9000 input power before using this procedure.
- If you are upgrading from an existing system with downtime that is included in the [supported systems matrix](#), you can upgrade controller hardware by moving storage or contact technical support. Refer to [References](#) to link to *Upgrade by moving volumes or storage*.

Automate the controller upgrade process

This procedure provides the steps for the automated procedure, which uses automatic disk assignment and network port reachability checks to simplify the controller upgrade experience.

Decide whether to use this aggregate relocation procedure

There are several aggregate relocation (ARL) methods for upgrading controller hardware. This content describes how to upgrade storage controllers in an HA pair by converting the existing system to the replacement system, keeping the existing system chassis and disks. You should only use this complex procedure if you're an experienced ONTAP administrator.

To help you decide if this ARL procedure is suitable for your controller hardware upgrade, you should review all of the following circumstances for supported and unsupported upgrades.

Supported for this ARL upgrade

You can use this ARL procedure under the following circumstances:

- Your controller upgrade is listed in the [supported systems matrix](#).
- You have verified with your NetApp sales representative that you have received the hardware that you need for the controller upgrade:
 - Two AFF A90, AFF A70, or AFF C80 controllers and all I/O modules. The required lengths of 100GbE cables.
 - Two AFF A50, AFF A30, AFF C30, or AFF C60 controllers and I/O modules and the required cables
 - ASA A150, AFF A150, or FAS2820 controller
 - ASA A900, AFF A900, or FAS9500 controller and NVRAM modules and parts
- You're running the minimum ONTAP version for your upgrade. For more information, refer to the [supported system upgrade combinations](#).
- You're experienced in administering ONTAP and are comfortable with the risks of working in diagnostic privilege mode.

- Your systems are running ONTAP 9.15.1 or later, and they are using Ethernet switches to connect to Ethernet-attached storage.



You can use NetApp Storage Encryption (NSE), NetApp Volume Encryption (NVE), and NetApp Aggregate Encryption (NAE) with this procedure.

Not supported for this ARL upgrade

You cannot use this ARL procedure under the following circumstances:

- You want to add the new controllers as a new HA pair to the cluster and migrate the data by using volume moves.
- You're upgrading a MetroCluster IP configuration.

To upgrade a MetroCluster IP configuration, see [References](#) to link to the *MetroCluster Upgrade and Expansion* content.

Supported system upgrade combinations

The following table shows the supported systems matrix for performing a controller upgrade by converting the existing system to the replacement system, keeping the existing system chassis and disks.



This procedure strictly applies to the following upgrade configurations. Don't use this procedure to perform an upgrade between any other system combinations. For all other controller models, refer to [References](#) to link to *Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later* and *Using aggregate relocation to manually upgrade controller hardware running ONTAP 9.8 or later*.

Existing system	Replacement system	Supported ONTAP versions
AFF C250 ¹	AFF C30, AFF C60	9.16.1
AFF A250 ¹	AFF A30, AFF A50	9.16.1
AFF C800 ¹	AFF C80	9.16.1
AFF A800 ¹	AFF A90 or AFF A70	9.15.1 and later
AFF A220 configured as an All SAN Array (ASA)	ASA A150	9.13.1P1 and later
AFF A220	AFF A150	9.10.1P15, 9.11.1P11, 9.12.1P5 and later
AFF A200	AFF A150	9.10.1P15, 9.11.1P11 and later ²
AFF C190	AFF A150	9.10.1P15, 9.11.1P11, 9.12.1P5 and later
FAS2620	FAS2820	9.11.1P7 or later patch releases (FAS2620) ² 9.13.1 and later (FAS2820)
FAS2720	FAS2820	9.13.1 and later
AFF A700 configured as an ASA	ASA A900	9.13.1P1 and later

Existing system	Replacement system	Supported ONTAP versions
AFF A700	AFF A900	9.10.1P10, 9.11.1P6 and later
FAS9000	FAS9500	9.10.1P10, 9.11.1P6 and later

¹ When you upgrade to a system introduced in ONTAP 9.15.1 or later, ONTAP converts the storage efficiency of all existing thin-provisioned volumes, including those not using storage efficiency, and applies the new storage efficiency features that make use of the hardware offload functionality. This is an automatic background process, with no visible performance impact to the system. [Learn more](#).

² The AFF A200 and FAS2620 systems don't support ONTAP versions later than 9.11.1.



NetApp strongly recommends, when possible, that you have the same ONTAP version on the old and replacement systems.

The minimum ONTAP versions in the preceding table are mandatory. These ONTAP versions have the Service Processor or baseboard management controller (BMC) firmware version that is required to support mixing controller types within a chassis during an upgrade.

Choose a different hardware upgrade procedure

- [Review the alternative ARL methods available for upgrading controller hardware](#).
- If you prefer a different method of upgrading the controller hardware and are willing to do volume moves, refer to [References](#) to link to *Upgrade by moving volumes or storage*.

Related information

Refer to [References](#) to link to the *ONTAP 9 Documentation*.

Required tools and documentation

You must have a grounding strap to perform the upgrade, and you need to reference other documents during the upgrade process.

For an AFF A800 upgrade to an AFF A90 or AFF A70, verify that the 100GbE cables are a minimum of one meter in length.

Refer to [References](#) to access the list of reference documents and reference sites required for this upgrade.

Guidelines for upgrading controllers

To understand whether you can use aggregate relocation (ARL), keeping the old system chassis and disks, depends on the system upgrade configuration and ONTAP version.

Supported upgrades for ARL

Controller upgrades are supported for certain system configurations. Refer to the [supported system upgrade combinations](#) for a list of the supported systems and minimum ONTAP version.

If you have received a new AFF A30, AFF A50, AFF A70, AFF A90, AFF A150, AFF A900, AFF C30, AFF C60, AFF C80, FAS2820, or FAS9500 as a complete system, including a new chassis, refer to [References](#) to link to *Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later*.

Controller upgrade using ARL is supported on systems configured with SnapLock Enterprise and SnapLock Compliance volumes.

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 don't need to convert them to a switched cluster.

Switch attached clusters

If you are upgrading nodes in a cluster that is connected to a cluster switch, you must verify that the make, model, firmware version, RCF, and ONTAP version running on the switch will be the same as those running on the replacement controller after the upgrade. If required, you must perform the switch upgrade before upgrading the controllers by using ARL.

For more information, see [Connect to a switch-attached cluster](#).

Troubleshoot

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.

If any problems occur, refer to the [Troubleshoot](#) section at the end of the procedure for more information and possible solutions. Information about the failures that can occur is listed by the phase of the procedure in the [ARL upgrade sequence](#).

If you don't find a solution to the problem you encountered, contact technical support.

Learn about the ARL upgrade sequence

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

Upgrade 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.

ARL upgrade sequence overview

During the procedure, you upgrade the original controller hardware with the replacement controller hardware, one controller at a time, taking advantage of the HA pair 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 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:

Stage	Steps
Stage 1: Prepare for upgrade	<p>During Stage 1, you verify that you have the correct hardware for your upgrade, run prechecks, and, if required, correct aggregate ownership. You must record certain information if you are managing Storage Encryption by using the Onboard Key Manager and you can choose to quiesce the SnapMirror relationships.</p> <p>Aggregate ownership at the end of Stage 1:</p> <ul style="list-style-type: none"> • Node1 is the home owner and current owner of the node1 aggregates • Node2 is the home owner and current owner of the node2 aggregates
Stage 2: Relocate resources and retire node1	<p>During Stage 2, you relocate node1 non-root aggregates and NAS data LIFs from node1 to node2. This process is largely automated; the operation pauses to enable you to check its status. You must manually resume the operation. If required, you relocate failed or vetoed aggregates. You record node1 information for use later in the procedure before retiring node1. You can also prepare to netboot node1 later in the procedure.</p> <p>Aggregate ownership at the end of Stage 2:</p> <ul style="list-style-type: none"> • Node2 is the current owner of node1 aggregates • Node2 is the home owner and current owner of node2 aggregates
Stage 3: Boot node1 with the replacement system modules	<p>During Stage 3, you boot node1 with upgraded system modules and verify the upgraded node1 installation. If you are using NetApp Volume Encryption (NVE), you restore key-manager configuration. You also relocate node1 non-root aggregates and NAS data LIFs from node2 to the upgraded node1 and verify that the SAN LIFs exist on node1.</p> <p>Aggregate ownership at the end of Stage 3:</p> <ul style="list-style-type: none"> • Upgraded node1 is the home owner and current owner of node1 aggregates • Node2 is the home owner and current owner of node2 aggregates
Stage 4: Relocate resources and retire node2	<p>During Stage 4, you relocate non-root aggregates and NAS data LIFs from node2 to the upgraded node1 and retire node2.</p> <p>Aggregate ownership at the end of Stage 4:</p> <ul style="list-style-type: none"> • Upgraded node1 is the home owner and current owner of aggregates that originally belonged to node1 • Upgraded node1 is the current owner of node2 aggregates

Stage	Steps
Stage 5: Install the replacement system modules on node2	<p>During Stage 5, you install the new system modules that you received for the upgraded node2 and then netboot node2.</p> <p>Aggregate ownership at the end of Stage 5:</p> <ul style="list-style-type: none"> Upgraded node1 is the home owner and current owner of the aggregates that originally belonged to node1. Upgraded node2 is the home owner and current owner of aggregates that originally belonged to node2.
Stage 6: Boot node2 with the replacement system modules	<p>During Stage 6, you boot node2 with upgraded system modules and verify the upgraded node2 installation. If you are using NVE, you restore key-manager configuration. For switch-attached cluster upgrades, you need to validate that the cluster switch reference configuration file (RCF) supports the shared cluster/HA ports. You also relocate node1 non-root aggregates and NAS data LIFs from node1 to the upgraded node2 and verify that the SAN LIFs exist on node2.</p>
Stage 7: Complete the upgrade	<p>During Stage 7, you confirm that the new nodes are set up correctly and, if the new nodes are encryption-enabled, you configure and set up Storage Encryption or NVE. You should also decommission the old nodes and resume the SnapMirror operations.</p>

Stage 1. Prepare for upgrade

Verify the upgrade hardware

Before starting the upgrade, verify that you have the correct modules for your replacement system. If there are parts missing, contact technical support or your NetApp sales representative for assistance.

If you're upgrading from ...	Replacement system must have ...
AFF A250, AFF C250	<ul style="list-style-type: none"> Two controller modules and new IO modules An X60132A, 4-port 10/25GbE card card for two-node switchless upgrade configurations
AFF A800, AFF C800	Two controller modules, two NVRAMs, and new IO modules
<ul style="list-style-type: none"> AFF A220 configured as an ASA AFF A220, AFF A200, AFF C190 FAS2620, FAS2720 	<p>Two controller modules</p> <p>If you're converting the existing system to a storage shelf so you can attach it to another system, the replacement system must also have two IO modules.</p>

If you're upgrading from ...	Replacement system must have ...
<ul style="list-style-type: none"> • AFF A700 configured as an ASA • AFF A700 • FAS9000 	Two controller and two NVRAM modules

Prepare the nodes for upgrade

The controller replacement process begins with a series of prechecks. You also gather information about the original nodes for use later in the procedure and, if required, determine the type of self-encrypting drives that are in use.

Steps

1. List the Service Processor (SP) or baseboard management controller (BMC) firmware version running on the old controller:

```
service-processor show
```

Verify that you have a supported SP or BMC firmware version:

Old controller	SP or BMC	Minimum firmware version
AFF A800	BMC	10.9
AFF A220	BMC	11.9P1
AFF A200	SP	5.11P1
AFF C190	BMC	11.9P1
FAS2620	SP	5.11P1
FAS2720	BMC	11.9P1

2. Begin the controller replacement process by entering the following command in the advanced privilege mode of the ONTAP command line:

```
set -privilege advanced
```

```
system controller replace start -nodes node_names
```

You will see output similar to the following example. The output displays the ONTAP version running on your cluster:

Warning:

1. Current ONTAP version is 9.15.1

2. Verify that NVMEM or NVRAM batteries of the new nodes are charged, and charge them if they are not. You need to physically check the new nodes to see if the NVMEM or NVRAM batteries are charged. You can check the battery status either by connecting to a serial console or using SSH, logging into the Service Processor (SP) or Baseboard Management Controller (BMC) 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.

3. If a controller was previously part of a different cluster, run wipeconfig before using it as the replacement controller.

4. Note: This is not a MetroCluster configuration. Controller replacement supports only ARL based procedures.

Do you want to continue? {y|n}: y

3. Select y. You will see the following output:

```
Controller replacement operation: Prechecks in progress.  
Controller replacement operation has been paused for user intervention.
```

During the prechecks phase, the system runs the following list of checks in the background.

Precheck	Description
Cluster Health Check	Checks all the nodes in the cluster to confirm that they are healthy.
Aggregate Relocation Status Check	Checks whether an aggregate relocation is already in progress. If another aggregate relocation is in progress, the check fails.
Model Name Check	Checks whether the controller models are supported for this procedure. If the models are not supported, the task fails.
Cluster Quorum Check	Checks that the nodes being replaced are in quorum. If the nodes are not in quorum, the task fails.

Precheck	Description
Image Version Check	Checks that the nodes being replaced run the same version of ONTAP. If the ONTAP image versions are different, the task fails. 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. For instructions on how to upgrade ONTAP, refer to References to link to <i>Upgrade ONTAP</i> .
HA Status Check	Checks if both the nodes being replaced are in a high availability (HA) pair configuration. If storage failover is not enabled for the controllers, the task fails.
Aggregate Status Check	If the nodes being replaced own aggregates for which they are not the home owner, the task fails. The nodes should not own any non-local aggregates.
Disk Status Check	If any nodes being replaced have missing or failed disks, the task fails. If any disks are missing, refer to References to link to <i>Disk and aggregate management with the CLI</i> , <i>Logical storage management with the CLI</i> , and <i>HA pair management</i> to configure storage for the HA pair.
Data LIF Status Check	Checks if any of the nodes being replaced have non-local data LIFs. The nodes should not contain any data LIFs for which they are not the home owner. If one of the nodes contains non-local data LIFs, the task fails.
Cluster LIF Status	Checks whether the cluster LIFs are up for both nodes. If the cluster LIFs are down, the task fails.
ASUP Status Check	If AutoSupport notifications are not configured, the task fails. You must enable AutoSupport before beginning the controller replacement procedure.
CPU Utilization Check	Checks if the CPU utilization is more than 50% for any of the nodes being replaced. If the CPU usage is more than 50% for a considerable period of time, the task fails.
Aggregate Reconstruction Check	Checks if reconstruction is occurring on any data aggregates. If aggregate reconstruction is in progress, the task fails.
Node Affinity Job Check	Checks if any node affinity jobs are running. If node affinity jobs are running, the check fails.

- After the controller replacement operation is started and the prechecks are completed, the operation pauses, enabling you to collect output information that you might need later in the controller upgrade process.
- Run the below set of commands as directed by the controller replacement procedure on the system console.

Run the commands from the serial port connected to each node, run and save the output of the commands individually:

- `vserver services name-service dns show`
- `network interface show -curr-node local -role cluster,intercluster,node-mgmt,cluster-mgmt,data`
- `network port show -node local -type physical`
- `service-processor show -node local -instance`
- `network fcp adapter show -node local`
- `network port ifgrp show -node local`
- `system node show -instance -node local`
- `run -node local sysconfig`
- `run -node local sysconfig -ac`
- `run -node local aggr status -r`
- `vol show -fields type`
- `run local aggr options data_aggregate_name`
- `vol show -fields type , space-guarantee`
- `storage aggregate show -node local`
- `volume show -node local`
- `storage array config show -switch switch_name`
- `system license show -owner local`
- `storage encryption disk show`
- `security key-manager onboard show-backup`
- `security key-manager external show`
- `security key-manager external show-status`
- `network port reachability show -detail -node local`



If NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE) using the Onboard Key Manager is in use, keep the key-manager passphrase ready to complete the key manager resync later in the procedure.

6. If your system uses self-encrypting drives, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use on the HA pair that you are upgrading. ONTAP software supports two types of self-encrypting drives:

- FIPS-certified NetApp Storage Encryption (NSE) SAS or NVMe drives
- Non-FIPS self-encrypting NVMe drives (SED)



You cannot mix FIPS drives with other types of drives on the same node or HA pair.

You can mix SEDs with non-encrypting drives on the same node or HA pair.

[Learn more about supported self-encrypting drives.](#)

Correct aggregate ownership if an ARL precheck fails

If the Aggregate Status Check fails, you must return aggregates owned by the partner node to the home owner node and initiate the precheck process again.

Steps

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

```
storage aggregate relocation start -node source_node -destination destination-  
node -aggregate-list *
```

2. 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
```

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.
```

After you finish

You must restart the controller replacement process:

```
system controller replace start -nodes node_names
```

License

Each node in the cluster must have its own NetApp License File (NLF).

If you do not have an NLF, 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 NLF for the new controller after the upgrade is complete.

Refer to [References](#) to link to the *NetApp Support Site* where you can obtain your NLF. The NLFs are available in the *My Support* section under *Software licenses*. If the site does not have the NLFs that you need, contact your NetApp sales representative.

For detailed information about licensing, refer to [References](#) to link to the *System Administration Reference*.

Manage storage encryption using the Onboard Key Manager

You can use the Onboard Key Manager (OKM) to manage encryption keys. If you have the OKM set up, you must record the passphrase and backup material before beginning the upgrade.

Steps

1. Record the cluster-wide passphrase.

This is the passphrase that was entered when the OKM was configured or updated using the CLI or REST API.

2. Back up the key-manager information by running the `security key-manager onboard show-backup` command.

Quiesce the SnapMirror relationships (optional)

Before continuing with the procedure, you must confirm that all the SnapMirror relationships are quiesced. When a SnapMirror relationship is quiesced, it remains quiesced across reboots and failovers.

Steps

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

```
snapmirror show
```



If the status is "Transferring", you must abort those transfers:

```
snapmirror abort -destination-vserver vs_server_name
```

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

2. Quiesce all relationships between the cluster:

```
snapmirror quiesce -destination-vserver *
```

Stage 2. Relocate resources and retire node1

Relocate non-root aggregates and NAS data LIFs owned by node1 to node2

Before you can replace node1 with the replacement modules for your system upgrade, you must move the non-root aggregates and NAS data LIFs from node1 to node2 before eventually restoring the node1 resources back on node1 running on the replacement system. This process is largely automated; the operation pauses to enable you to check its status.

Before you begin

The operation should already be paused when you begin the task; you must manually resume the operation.

About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. You are not required to move SAN LIFs for cluster or service health during the upgrade. You must verify that the LIFs are healthy and located on

appropriate ports after you bring node1 online as the replacement system.



The home owner for the aggregates and LIFs is not modified; only the current owner is modified.

Steps

1. Resume the aggregate relocation and NAS data LIF move operations:

```
system controller replace resume
```

All the non-root aggregates and NAS data LIFs are migrated from node1 to node2.

The operation pauses to enable you to verify whether all node1 non-root aggregates and non-SAN data LIFs have been migrated to node2.

2. Check the status of the aggregate relocation and NAS data LIF move operations:

```
system controller replace show-details
```

3. With the operation still paused, verify that all the non-root aggregates are online for their state on node2:

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

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,normal								
aggr_2	825.0GB	825.0GB	0%	online	1	node2		
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 <aggregate_name>
```

4. Verify that all the volumes are online on node2 by using 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 <vserver_name> -volume <volume_name>
```

The `vserver_name` to use with this command is found in the output of the previous `volume show` command.

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

```
network interface modify -vserver vserver_name -lif LIF_name -home-node  
nodename -status-admin up
```

Relocate failed or vetoed aggregates to node2

If any aggregates fail to relocate or are vetoed, you need to manually relocate the aggregates to node2, or if necessary, override either the vetoes or destination checks.

About this task

The system pauses the relocation operation due to the error.

Steps

1. Check the event management system (EMS) logs to determine why the aggregate failed to relocate or was vetoed.
2. Relocate any failed or vetoed aggregates:

```
storage aggregate relocation start -node <node1> -destination <node2>  
-aggregate-list <aggregate_name> -ndo-controller-upgrade true
```

3. When prompted, enter `y`.
4. You can force relocation by using one of the following methods:

Option	Description
Overriding veto checks	Use the following command: <pre>storage aggregate relocation start -node node1 -destination node2 -aggregate-list <aggregate_list> -ndo-controller-upgrade true -override-vetoes true</pre>
Overriding destination checks	Use the following command: <pre>storage aggregate relocation start -node node1 -destination node2 -aggregate-list <aggregate_list> -ndo-controller-upgrade true -override-vetoes true -override-destination-checks true</pre>

Retire node1

To retire node1, you resume the automated operation to disable the HA pair with node2 and shut down node1 correctly.

Steps

1. Resume the operation:

```
system controller replace resume
```

2. Verify that node1 has been halted:

```
system controller replace show-details
```

After node1 has completely halted, node1 should be at the LOADER> prompt. To see the LOADER> prompt, connect to the serial console of node1.

Replace the node1 system modules

Replace the AFF A250 or AFF C250 controller modules

At this stage, node1 is down and all data is served by node2. You must take care to remove only the node1 controller module. Typically, node1 is controller A, located on the left side of the chassis when looking at the controllers from the rear of the system. The controller label is located on the chassis directly above the controller module.



Don't power off the chassis because node1 and node2 are in the same chassis and connected to the same power supplies.

Remove the AFF A250 or AFF C250 controller module

To remove the node1 controller module, you first remove the cable management device, unlock the locking latches, and then remove the controller module from the chassis.

Before you begin


If you are not already grounded, correctly ground yourself.

Steps

1. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.



If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).



1	Lever
2	Latching mechanism

2. Go to the rear of the chassis.
3. Unplug the node1 controller module power supply from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies for node1.



The power connections for node1 and node2 are on top of each other. Take care to only unplug the cables for node1. Unplugging the cables for node1 and node2 could cause a power outage to both nodes in the HA pair.

5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

8. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.

Make sure that you support the weight of the controller module as you slide it out of the chassis.

Install the AFF A30, AFF A50, AFF C30, or AFF C60 controller module

Install, cable, and connect the replacement module in node1.

Before you begin

Verify that you have an X60132A, 4-port 10/25 GbE card in slot1 on node1. The X60132A card is required for cluster interconnect on two-node switchless cluster configurations during the upgrade.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.



Don't completely insert the controller module in the chassis until instructed to do so later in the procedure.

2. Cable the management and console ports to the node1 controller module.



Because the chassis is already powered ON, node1 starts BIOS initialization followed by AUTOBOOT as soon as you insert the new controller module. To avoid this AUTOBOOT, NetApp recommends connecting the serial and console cables before inserting the controller module.

3. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated. The locking latch rises when the controller module is fully seated. Close the cam handle to the locked position.



To avoid damaging the connectors, don't use excessive force when sliding the controller module into the chassis.

4. Connect the serial console as soon as the module is seated and be ready to interrupt AUTOBOOT of node1.
5. After you interrupt AUTOBOOT, node1 stops at the LOADER prompt.

If you do not interrupt AUTOBOOT on time and node1 starts booting, wait for the prompt and press Ctrl-C to go into the boot menu. After the node stops at the boot menu, use option 8 to reboot the node and interrupt AUTOBOOT during reboot.

6. At the LOADER> prompt of node1, set the default environment variables:

```
set-defaults
```

7. Save the default environment variables settings:

```
saveenv
```

Replace the AFF A800 or AFF C800 controller modules

At this stage, node1 is down and all data is served by node2. You must take care to remove only the node1 controller module. Typically, node1 is controller A, located on the left side of the chassis when looking at the controllers from the rear of the system. The controller label is located on the chassis directly above the controller module.



Don't power off the chassis because node1 and node2 are in the same chassis and connected to the same power supplies.

Before you begin

If you are not already grounded, correctly ground yourself.


Remove the AFF A800 or AFF C800 controller module

Remove the cable management device from the existing module and move the controller slightly out of the chassis.

Steps

1. Prepare to remove the controller module:
 - a. On the front of the chassis, use your thumbs to firmly push each drive in to the top and bottom disk bays until you feel a positive stop. This ensures that the drives are firmly seated against the chassis

midplane.




1. Unplug the node1 controller module power supplies from the source.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.



1	Locking latch
2	Locking pin

Install the AFF A90, AFF A70, or AFF C80 controller module


Install, cable, and connect the AFF A90, AFF A70, or AFF C80 controller module in node1.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.



Don't completely insert the controller module in the chassis until instructed to do so later in the procedure.




2. Cable the management and console ports to the node1 controller module.



Because the chassis is already powered ON, node1 starts BIOS initialization followed by AUTOBOOT as soon as you insert the new controller module. To avoid this AUTOBOOT, NetApp recommends connecting the serial and console cables before inserting the controller module.

3. On the front of the chassis, use your thumbs to firmly push each drive in to the top and bottom disk bays until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.



- a. Go to the rear of the chassis.
4. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated. The locking latch rises when the controller module is fully seated. Close the cam handle to the locked position.



To avoid damaging the connectors, don't use excessive force when sliding the controller module into the chassis.

5. Connect the serial console as soon as the module is seated and be ready to interrupt AUTOBOOT of node1.
6. After you interrupt AUTOBOOT, node1 stops at the LOADER prompt.

If you do not interrupt AUTOBOOT on time and node1 starts booting, wait for the prompt and press Ctrl-C to go into the boot menu. After the node stops at the boot menu, use option 8 to reboot the node and interrupt AUTOBOOT during reboot.

7. At the LOADER> prompt of node1, set the default environment variables:

```
set-defaults
```

8. Save the default environment variables settings:

```
saveenv
```

Replace the AFF A220, AFF A200, AFF C190, FAS2620, or FAS2720 controller module

At this stage, node1 is down and all data is served by node2. You must take care to remove only the node1 controller module. Typically, node1 is controller A, located on the left side of the chassis when looking at the controllers from the rear of the system. The controller label is located on the chassis directly above the controller module.



Don't power off the chassis because node1 and node2 are in the same chassis and connected to the same power supplies.

Before you begin

If you are not already grounded, correctly ground yourself.

Remove the AFF A220, AFF A200, AFF C190, FAS2620, or FAS2720 controller module


To access components inside the controller, remove the controller module from the system and then remove the cover on the controller module.

Steps


1. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

2. Remove and set aside the cable management devices from the left and right sides of the controller module.



3. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.



4. Turn the controller module over and place it on a flat, stable surface.

Install the ASA A150, AFF A150, or FAS2820 controller module

Install, cable, and connect the ASA A150, AFF A150, or FAS2820 controller module in node1.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.



Don't completely insert the controller module in the chassis until instructed to do so later in the procedure.

2. Cable the management and console ports to the node1 controller module.



Because the chassis is already powered ON, node1 starts BIOS initialization followed by AUTOBOOT as soon as it is fully seated. To interrupt the node1 boot, before completely inserting the controller module into the slot, it is recommended that you connect the serial console and management cables to the node1 controller module.

3. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated. The locking latch rises when the controller module is fully seated. Close the cam handle to the locked position.



To avoid damaging the connectors, don't use excessive force when sliding the controller module into the chassis.

4. Connect the serial console as soon as the module is seated and be ready to interrupt AUTOBOOT of node1.
5. After you interrupt AUTOBOOT, node1 stops at the LOADER prompt. If you do not interrupt AUTOBOOT on time and node1 starts booting, wait for the prompt and press Ctrl-C to go into the boot menu. After the node stops at the boot menu, use option 8 to reboot the node and interrupt AUTOBOOT during reboot.
6. At the LOADER> prompt of node1, set the default environment variables:

```
set-defaults
```

7. Save the default environment variables settings:

```
saveenv
```

Replace the AFF A700 or FAS9000 controller and NVRAM modules

At this stage, node1 is down and all data is served by node2. You must take care to remove only the node1 controller module and the node1 NVRAM module. Typically, node1 is controller A, located on the left side of the chassis when looking at the controllers from the rear of the system. The controller label is located on the chassis directly above the controller module.



Don't power off the chassis because node1 and node2 are in the same chassis and connected to the same power supplies.

Before you begin

If you are not already grounded, correctly ground yourself.

Remove the AFF A700 or FAS9000 controller module

Detach and remove the AFF A700 or FAS9000 controller module from node1.


Steps

1. Detach the console cable, if any, and the management cable from the node1 controller module.



When you are working on node1, you only remove the console and e0M cables from node1. You must not remove or change any other cables or connections on either node1 or node2 during this process.

2. Unlock and remove the controller module A from the chassis.
 - a. Slide the orange button on the cam handle downward until it unlocks.



1	Cam handle release button
2	Cam handle

- b. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Remove the AFF A700 or FAS9000 NVRAM module

Unlock and remove the AFF A700 or FAS9000 NVRAM module from node1.



The AFF A700 or FAS9000 NVRAM module is in slot 6 and is double the height of the other modules in the system.

Steps

1. Unlock and remove the NVRAM module from slot 6 of node1.


- a. Depress the lettered and numbered cam button.

The cam button moves away from the chassis.

- b. Rotate the cam latch down until it is in a horizontal position.

The NVRAM module disengages from the chassis and moves a few inches.

- c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module face.



1	Lettered and numbered I/O cam latch
2	I/O latch completely unlocked

Install the ASA A900, AFF A900, or FAS9500 NVRAM and controller modules

Install, cable, and connect the ASA A900, AFF A900, or FAS9500 NVRAM and controller modules in node1.

You must note the following when performing the installation:

- Move all blank filler modules in slots 6-1 and 6-2 from the old NVRAM module to the new NVRAM module.
- Do NOT move the coredump device from the AFF A700 NVRAM module to the ASA A900 or AFF A900 NVRAM module.
- Move all flash cache modules installed in the FAS9000 NVRAM module to the FAS9500 NVRAM module.

Before you begin


If you are not already grounded, correctly ground yourself.

Install the ASA A900, AFF A900, or FAS9500 NVRAM module

Install the ASA A900, AFF A900, or FAS9500 NVRAM module in slot 6 of node1.

Steps

1. Align the NVRAM module with the edges of the chassis opening in slot 6.
2. Gently slide the NVRAM module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the NVRAM module in place.



1	Lettered and numbered I/O cam latch
2	I/O latch completely unlocked

Install the ASA A900, AFF A900, or FAS9500 controller module on node1.

Use the following procedure to install the ASA A900, AFA A900, or FAS9500 controller module in node1.

Steps

1. Align the end of the controller module with opening A in the chassis, and then gently push the controller module halfway into the system.



Don't completely insert the controller module in the chassis until instructed to do so later in the procedure.

2. Cable the management and console ports to the node1 controller module.




Because the chassis is already powered ON, node1 starts BIOS initialization followed by AUTOBOOT as soon as it is fully seated. To interrupt the node1 boot, before completely inserting the controller module into the slot, it is recommended that you connect the serial console and management cables to the node1 controller module.

3. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latch rises when the controller module is fully seated.



To avoid damaging the connectors, don't use excessive force when sliding the controller module into the chassis.



1	Cam handle locking latch
2	Cam handle in the unlocked position

4. Connect the serial console as soon as the module is seated and be ready to interrupt AUTOBOOT of node1.
5. After you interrupt AUTOBOOT, node1 stops at the LOADER prompt. If you do not interrupt AUTOBOOT on time and node1 starts booting, wait for the prompt and press Ctrl-C to go into the boot menu. After the node stops at the boot menu, use option 8 to reboot the node and interrupt AUTOBOOT during reboot.
6. At the LOADER> prompt of node1, set the default environment variables:

```
set-defaults
```

7. Save the default environment variables settings:

```
saveenv
```

Netboot node1

After swapping the corresponding replacement system modules, you must netboot node1. The term netboot means that you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you add a copy of the ONTAP 9 boot image onto a web server that the system can access.

It isn't possible to check the version of ONTAP installed on the boot media of the replacement controller module unless it is installed in a chassis and powered ON. The ONTAP version on the replacement system boot media must be same as the ONTAP version running on the old system that you are upgrading and both the primary and backup boot images on the boot media must match. To verify the minimum supported ONTAP

version for your upgrade, see the [supported systems matrix](#).

You can configure the images by performing a netboot followed by the `wipeconfig` command from the boot menu. If the controller module was previously used in another cluster, the `wipeconfig` command clears any residual configuration on the boot media.

You can also use the USB boot option to perform the netboot. See the Knowledge Base article [How to use the boot_recovery LOADER command for installing ONTAP for initial setup of a system](#).

Before you begin

- Verify that you can access a HTTP server with the system.
- Download the necessary system files for your system and the correct version of ONTAP from the *NetApp Support Site*. Refer to [References](#) to link to the *NetApp Support Site*.

About this task

You must netboot the new controllers if they do not have the same 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.


Steps

1. Refer to [References](#) to link to the *NetApp Support Site* to download the files used for performing the netboot of the system.
2. Download the appropriate ONTAP software from the software download section of the *NetApp Support Site* and store the `<ontap_version>_image.tgz` file on a web-accessible directory.
3. Change to the web-accessible directory and verify that the files you need are available.
4. Your directory listing should contain `<ontap_version>_image.tgz`.
5. Configure the netboot connection by choosing one of the following actions.



You must 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 using the following command at the boot environment prompt: <code>ifconfig e0M -auto</code>

If Dynamic Host Configuration Protocol (DHCP) is...	Then...
Not running	<p>Manually configure the connection by using the following command at the boot environment prompt:</p> <pre>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></pre> <p><i>filer_addr</i> is the IP address of the storage system (mandatory). <i>netmask</i> is the network mask of the storage system (mandatory). <i>gateway</i> is the gateway for the storage system (mandatory). <i>dns_addr</i> is the IP address of a name server on your network (optional). <i>dns_domain</i> is the Domain Name Service (DNS) domain name (optional).</p> <div>  <p>Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p> </div>

6. Perform netboot on node1:

```
netboot http://<web_server_ip/path_to_web_accessible_directory>/netboot/kernel
```



Do not interrupt the boot.

7. (AFF A250 and AFF C250 upgrades only) When node1 for the replacement controller module is booting, the following warning displays because the configuration includes node2 for the existing controller:

```
*****
* WARNING: Partner is not of the same family/model. *
* Mixing is only allowed when upgrading the system. *
* The system will shut down in 24 hours.             *
*****
Do you want to continue (y/n):
```

Answer `y`.

This warning displays for every system boot until you upgrade node2. This is the expected behavior.

8. Wait for the node1 running on the replacement controller module to boot and display the boot menu options as shown below:

Please choose one of the following:

- (1) Normal Boot.
 - (2) Boot without /etc/rc.
 - (3) Change password.
 - (4) Clean configuration and initialize all disks.
 - (5) Maintenance mode boot.
 - (6) Update flash from backup config.
 - (7) Install new software first.
 - (8) Reboot node.
 - (9) Configure Advanced Drive Partitioning.
 - (10) Set Onboard Key Manager recovery secrets.
 - (11) Configure node for external key management.
- Selection (1-11)?

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

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

Disregard the following message:

This procedure is not supported for Non-Disruptive Upgrade on an HA pair

This note applies to nondisruptive ONTAP software upgrades, and not controller upgrades.



Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the incorrect image might install. This issue applies to all ONTAP releases. The netboot procedure combined with option (7) Install new software wipes the boot media and places the same ONTAP version on both image partitions.

10. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL:

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

The `<path_to_the_web-accessible_directory>` should lead to where you downloaded the `<ontap_version>_image.tgz` in [Step 2](#).

11. Complete the following substeps to reboot the controller module:

- 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}
```

- b. Enter `y` to reboot 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 must be restored.

12. Clear any previous configuration on the boot media.

- a. At the following prompt, run the `wipeconfig` command, and press the enter key:

```
Please choose one of the following:

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? wipeconfig
```

- b. When you see the message below, answer `yes`:

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

- c. The node reboots to finish the `wipeconfig` and then stops at the boot menu.



Wait until the node stops at the boot menu after completing the `wipeconfig` operation.

13. Select option 5 to go to maintenance mode from the boot menu. Answer `yes` to the prompts until the node stops at maintenance mode and the command prompt `*>`.
14. Verify that the controller and chassis are configured as `ha`:

```
ha-config show
```

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

```
Chassis HA configuration: ha
Controller HA configuration: ha
```

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

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

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

16. Verify the `ha-config` settings:

```
ha-config show
```

```
Chassis HA configuration: ha
Controller HA configuration: ha
```

17. Halt node1:

```
halt
```

Node1 should stop at the LOADER prompt.

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

```
date
```

19. On node1, check the date by using the following command at the boot environment prompt:

```
show date
```

20. If necessary, set the date on node1:

```
set date mm/dd/yyyy
```



Set the corresponding UTC date on node1.

21. On node1, check the time by using the following command at the boot environment prompt:

```
show time
```

22. If necessary, set the time on node1:

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



Set the corresponding UTC time on node1.

23. Set the partner system ID on node1:


```
setenv partner-sysid node2_sysid
```

For node1, the `partner-sysid` must be that of node2. You can obtain the node2 system ID from the `node show -node node2` command output on node2.

- a. Save the settings:

```
saveenv
```

24. On node1, at the LOADER prompt, verify the `partner-sysid` for node1:

```
printenv partner-sysid
```

Stage 3. Boot node1 with the replacement system modules

Cable node1 for shared cluster-HA and storage

If you are performing one of the following upgrades, you must connect the cluster, HA, storage, data, and management connections that were previously connected to the node1 on the existing system to the newly installed node1 on the replacement system.

Existing system	Replacement system
AFF A250	AFF A30, AFF A50
AFF C250	AFF C30, AFF C60
AFF A800	AFF A70, AFF A90
AFF C800	AFF C80

Connect the e0M and BMC ports

If the existing system has a management port (e0M) and a BMC port, the e0M and BMC ports are combined and accessed through the "wrench" port on the replacement system. You must ensure that the e0M and BMC ports are connected to the same switch and subnet on the existing system before connecting to the replacement system.

If the...	Then...
e0M and BMC IP addresses are on the same IP subnet	Connect either the e0M or BMC port on the existing system to the "wrench" port on the replacement system.
e0M and BMC IP addresses are on different subnets	<ol style="list-style-type: none">1. Merge the e0M and BMC IP addresses into one IP subnet.2. Connect either the e0M or BMC port on the existing system to the "wrench" port on the replacement system.

Connect to a two-node switchless cluster

The following tables show the switch port usage for two-node switchless cluster configurations.

Port type	AFF A800, AFF C800	AFF A90	AFF A70, AFF C80
Cluster	e0a	e1a	e1a
Cluster	e1a	e7a (Use e1b if there is no e7a)	e1b
HA	e0b	Don't connect	Don't connect
HA	e1b	Don't connect	Don't connect
SAS storage ports (if present and used)	Any available port	Any available port	Any available port
Ethernet storage ports for NS224 shelves	Any available port	Refer to Ethernet storage connectivity mapping	Refer to Ethernet storage connectivity mapping

Port	AFF A250, AFF C250	AFF A30, AFF C30, AFF C60	AFF A50
Cluster	e0c	e1a (Use e1a for temporary cluster interconnect)	e1a (Use e1a for temporary cluster interconnect)
Cluster	e0d	e1b (Use e1b for temporary cluster interconnect)	e1b (Use e1b for temporary cluster interconnect)
HA	Not required	HA ports aren't required for the node1 upgrade	HA ports aren't required for the node1 upgrade
Ethernet storage Ports	Any available port	e3a, e3b	e3a, e3b
SAS storage ports	Any available port	3a, 3b	3a, 3b

Connect to a switch-attached cluster

For a switch-attached cluster, check that you meet the following requirements for the AFF A30, AFF A50, AFF A70, AFF A90, AFF C30, AFF C60, or AFF C80 (replacement) node:

- The identical cluster ports on the replacement node are on the same switch. For example, on completion of the upgrade, e1a on node1 and e1a on node2 should be attached to one cluster switch. Similarly, the second cluster port from both nodes should be attached to the second cluster switch. Cross-connection of shared cluster-HA ports, where e1a from node1 is connected to switchA and e1a from node2 is connected to switchB, results in HA communication failures.
- The replacement node uses shared cluster-HA Ethernet ports.
- Verify that the cluster switches are installed with a reference configuration file (RCF) that supports shared cluster-HA ports:
 1. Remove the existing configuration on the switch:

If your switch model is...	Go to...
Cisco Nexus	The Knowledge Base article How to clear configuration on a Cisco interconnect switch while retaining remote connectivity
Broadcom BES-53248	The Knowledge Base article How to clear configuration on a Broadcom interconnect switch while retaining remote connectivity

2. Configure and verify the switch setup:

If your switch model is...	Go to...
Cisco Nexus 9336C-FX2	Upgrade your Reference Configuration File (RCF)
Broadcom BES-53248	Upgrade the Reference Configuration File (RCF)
NVIDIA SN2100	Install or upgrade the Reference Configuration File (RCF) script



If the cluster switch only supports 10/25 GbE speeds, you must use an X60130A, 4-port 10/25GbE card in slot1 or slot2 on the replacement system for cluster interconnect.

Boot node1 with the replacement system modules

Node1 with the replacement modules is now ready to boot. The supported replacement modules are listed in the [supported systems matrix](#).



When replacing controller modules, move all connections from the old to the replacement controller module.

When replacing the controller and NVRAM modules, move only the console and management connections.

Steps

1. (AFF A250, AFF C250, AFF A800, or AFF C800 upgrade only) At the LOADER prompt, enter maintenance mode:

```
boot_ontap maint
```

- a. Answer `y` to the mixed platform confirmation prompt.
- b. Answer `yes` to the confirmation prompt.
- c. Show the state of the 100GbE interfaces:

```
storage port show.
```

All 100GbE ports connected to NS224 shelves or storage switches should report as `storage` ports, as shown in the example output below.

```
*> storage port show
Port Type Mode      Speed (Gb/s) State      Status  VLAN ID
----
e8a  ENET storage 100 Gb/s    enabled  online  30
e8b  ENET storage 100 Gb/s    enabled  online  30
e11a ENET storage 100 Gb/s    enabled  online  30
e11b ENET storage 100 Gb/s    enabled  online  30
```

d. Exit maintenance mode:

```
halt
```

2. If you have NetApp Storage Encryption (NSE) drives installed, perform the following steps:



If you have not already done so earlier in the procedure, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use.

a. Set `bootarg.storageencryption.support` to `true` or `false`:

If the following drives are in use...	Then...
NSE drives that conform to FIPS 140-2 Level 2 self-encryption requirements	<code>setenv bootarg.storageencryption.support true</code>
NetApp non-FIPS SEDs	<code>setenv bootarg.storageencryption.support false</code>



You cannot mix FIPS drives with other types of drives on the same node or HA pair. You can mix SEDs with non-encrypting drives on the same node or HA pair.

b. Go to the special boot menu and select option (10) `Set Onboard Key Manager recovery secrets`.

Enter the passphrase and the backup information that you recorded earlier procedure. See [Manage storage encryption using the Onboard Key Manager](#).

3. Boot the node into the boot menu:

```
boot_ontap menu
```

4. When the node stops at the boot menu, reassign the old node1 disks to the replacement node1 by running the following command on node1:

```
boot_after_controller_replacement
```

After a short delay, you are prompted to enter the name of the node that is being replaced. If there are shared disks (also called Advanced Disk Partitioning (ADP) or partitioned disks), you are prompted to enter the node name of the HA partner.

These prompts might get buried in the console messages. If you do not enter a node name or enter an incorrect name, you are prompted to enter the name again.

If [localhost:disk.encryptNoSupport:ALERT]: Detected FIPS-certified encrypting drive and, or, [localhost:diskown.errorDuringIO:error]: error 3 (disk failed) on disk errors occur, perform the following steps:



1. Halt the node at the LOADER prompt.
2. Check and reset the storage encryption bootargs mentioned in [Step 2](#).
3. At the LOADER prompt, boot up:

```
boot_ontap
```

You can use the following example as a reference:

Expand the console output example

```
LOADER-A> boot_ontap menu
.
.
<output truncated>
.
All rights reserved.
*****
*
* Press Ctrl-C for Boot Menu. *
*
*****
.
<output truncated>
.
Please choose one of the following:

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 22/7

(22/7) Print this secret List
(25/6) Force boot with multiple filesystem
disks missing.
(25/7) Boot w/ disk labels forced to clean.
(29/7) Bypass media errors.
(44/4a) Zero disks if needed and create new
flexible root volume.
(44/7) Assign all disks, Initialize all
disks as SPARE, write DDR labels
.
.
<output truncated>
.
.
(wipeconfig) Clean all configuration on boot
```

```

device
(boot_after_controller_replacement) Boot after controller upgrade
(boot_after_mcc_transition)          Boot after MCC transition
(9a)                                Unpartition all disks and remove
their ownership information.
(9b)                                Clean configuration and
initialize node with partitioned disks.
(9c)                                Clean configuration and
initialize node with whole disks.
(9d)                                Reboot the node.
(9e)                                Return to main boot menu.

```

The boot device has changed. System configuration information could be lost. Use option (6) to restore the system configuration, or option (4) to initialize all disks and setup a new system. Normal Boot is prohibited.

Please choose one of the following:

```

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? boot_after_controller_replacement

```

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

```

.
.
<output truncated>
.
.
Controller Replacement: Provide name of the node you would like to
replace:<nodename of the node being replaced>
Controller Replacement: Provide High Availability partner of node1:
<nodename of the partner of the node being replaced>

```

```

Changing sysid of node nodel disks.
Fetched sanown old_owner_sysid = 536940063 and calculated old sys id
= 536940063
Partner sysid = 4294967295, owner sysid = 536940063
.
.
<output truncated>
.
.
varfs_backup_restore: restore using /mroot/etc/varfs.tgz
varfs_backup_restore: attempting to restore /var/kmip to the boot
device
varfs_backup_restore: failed to restore /var/kmip to the boot device
varfs_backup_restore: attempting to restore env file to the boot
device
varfs_backup_restore: successfully restored env file to the boot
device wrote key file "/tmp/rndc.key"
varfs_backup_restore: timeout waiting for login
varfs_backup_restore: Rebooting to load the new varfs
Terminated
<node reboots>

System rebooting...

.
.
Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...
Rebooting to load the restored env file...
.
System rebooting...
.
.
.
<output truncated>
.
.
.
.
WARNING: System ID mismatch. This usually occurs when replacing a
boot device or NVRAM cards!
Override system ID? {y|n} y
.
.
.

```


.
Login:



The system IDs shown in the preceding example are example IDs. The actual system IDs of the nodes that you are upgrading will be different.

Between entering node names at the prompt and the login prompt, the node reboots a few times to restore the environment variables, update firmware on the cards in the system, and for other ONTAP updates.

Restore key-manager configuration on the upgraded node1

If you are using NetApp Aggregate Encryption (NAE) or NetApp Volume Encryption (NVE) to encrypt volumes on the system you are upgrading, the encryption configuration must be synchronized to the new nodes. If you do not resynchronize the key-manager, when you relocate the node1 aggregates from node2 to the upgraded node1 by using ARL, failures might occur because node1 does not have the required encryption keys to bring encrypted volumes and aggregates online.

About this task

Synchronize the encryption configuration to the new nodes by performing the following steps:

Steps

1. Run the following command from node1:

```
security key-manager onboard sync
```

2. Verify that the SVM-KEK key is restored to "true" on node1 before you relocate the data aggregates:

```
::> security key-manager key query -node node1 -fields restored -key  
-type SVM-KEK
```

Example

```
::> security key-manager key query -node node1 -fields restored -key  
-type SVM-KEK
```

node	vserver	key-server	key-id
restored			
-----	-----	-----	-----
node1	svm1	""	0000000000000000020000000000a008a81976
true			2190178f9350e071fbb90f00000000000000000

Move node1 non-root aggregates and NAS data LIFs from node2 to the upgraded node1

After verifying the network configuration on node1, you need to relocate the NAS data LIFs owned by node1 from node2 to node1 and confirm that the SAN LIFs exist on node1.

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 the correct ports after you bring node1 online.

Steps

- 1. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

- Cluster quorum check
- System ID check
- Image version check
- Target platform check
- Network reachability check

The system pauses the operation at this stage in the network reachability check

- 2. Perform a network reachability check:

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

Confirm that all connected ports, including the interface group and VLAN ports, show their status as OK.

- 3. For the following upgrades, you must reassign the FCP SAN LIFs.

Existing system	Replacement system
AFF A250	AFF A30, AFF A50
AFF C250	AFF C30, AFF C60
AFF A800	AFF A70, AFF A90
AFF C800	AFF C80

For all other system upgrades, proceed to [Step 4](#).

- a. Reassign FCP SAN LIFs used for FCP or FC-NVMe data access to the correct home ports:

```
network interface show -vserver <vserver_hosting_fcp_lifs>
```

- b. For LIFs with the current node as the upgraded node1 and the current port reports "status oper" as "-" (because the port existed on the AFF A800 node but does not exist on the AFF A90 node), modify the current port before it can be brought online.

Verify that physical connectivity is established to the FC target port where the FC LIF needs to be moved:

- i. Set the LIF status to "down":

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

- ii. Modify the home port of the LIF:

```
network interface modify -vserver <vserver_name> -lif <lif_name> -
home-node <node1> -home-port <FC_target_port>
```

- iii. Set the LIF status to "up":

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

Repeat Substeps a and b for each FC SAN LIF that is home on node1.

4. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following checks:

- Cluster health check
- Cluster LIF status check

After performing these checks, the system relocates the non-root aggregates and NAS data LIFs owned by node1 to the new node1.

The controller replacement operation pauses after the resource relocation is complete.

5. Check the status of the aggregate relocation and NAS data LIF move operations:

```
system controller replace show-details
```

If the controller replacement procedure is paused, check and correct the error, if any, and then issue `resume` to continue the operation.

6. If necessary, restore and revert displaced LIFs or manually migrate and modify the node1 LIFs that failed to relocate automatically to node1.

Restore and revert displaced LIFs

- a. List any displaced LIFs:

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

- b. If any LIFs are displaced, restore the home node back to node1:

```
cluster controller-replacement network displaced-interface  
restore-home-node -node <node1_nodename> -vserver <vserver name>  
-lif-name <lif_name>
```

Manually migrate and modify LIFs

- a. Migrate the LIFs that failed to relocate automatically to node1:

```
network interface migrate -vserver <vserver name> -lif <lif_name>  
-destination-node <node1_nodename> -destination-port  
<port_on_node1>
```

- b. Modify the home node and home port for the migrated LIFs:

```
network interface modify -vserver <vserver_name> -lif  
<data_lif_name> -home-node <node1_nodename> -home-port  
<home_port>
```

7. Resume the operation to prompt the system to perform the required post-checks:

```
system controller replace resume
```

The system performs the following post-checks:

- Cluster quorum check
- Cluster health check
- Aggregates reconstruction check
- Aggregate status check

- Disk status check
- Cluster LIF status check
- Volume check

Stage 4. Relocate resources and retire node2

Relocate non-root aggregates and NAS data LIFs from node2 to node1

Before you can replace node2 with the replacement system module, you must first relocate the non-root aggregates that are owned by node2 to node1.

Before you begin

After the post-checks from the previous stage complete, the resource release for node2 starts automatically. The non-root aggregates and non-SAN data LIFs are migrated from node2 to the new node1.

About this task

After the aggregates and LIFs are migrated, the operation is paused for verification purposes. At this stage, you must verify that all the non-root aggregates and non-SAN data LIFs are migrated to the new node1.

The home owner for the aggregates and LIFs are not modified; only the current owner is modified.

Steps

1. Verify that all the non-root aggregates are online and their state on node1:

```
storage aggregate show -node node1 -state online -root false
```

The following example shows that the non-root aggregates on node1 are online:

```
cluster::> storage aggregate show -node node1 state online -root false
```

Aggregate	Size	Available	Used%	State	#Vols	Nodes
RAID	Status					
-----	-----	-----	-----	-----	-----	-----
-----	-----					
aggr_1	744.9GB	744.8GB	0%	online	5	node1
raid_dp	normal					
aggr_2	825.0GB	825.0GB	0%	online	1	node1
raid_dp	normal					
2 entries were displayed.						

If the aggregates have gone offline or become foreign on node1, bring them online by using the following command on the new node1, once for each aggregate:

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

2. Verify that all the volumes are online on node1 by using the following command on node1 and examining its output:

```
volume show -node node1 -state offline
```

If any volumes are offline on node1, bring them online by using the following command on node1, 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.

3. Verify that the LIFs have been moved to the correct ports and have a status of up. 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
```

4. Verify that there are no data LIFs remaining on node2 by using the following command and examining the output:

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

Relocate failed or vetoed aggregates to node1

If any aggregates fail to relocate or are vetoed, you need to manually relocate the aggregates to node1, or if necessary, override either the vetoes or destination checks.

About this task

The system pauses the relocation operation due to the error.

Steps

1. Check the event management system (EMS) logs to determine why the aggregate failed to relocate or was vetoed.
2. Relocate any failed or vetoed aggregates:

```
storage aggregate relocation start -node <node2> -destination <node1>  
-aggregate-list <aggregate_name> -ndo-controller-upgrade true
```

3. When prompted, enter *y*.
4. You can force relocation by using one of the following methods:

Option	Description
Overriding veto checks	Use the following command: storage aggregate relocation start -node node2 -destination node1 -aggregate-list <aggregate_list> -ndo-controller-upgrade true -override-vetoes true

Option	Description
Overriding destination checks	Use the following command: <pre>storage aggregate relocation start -node node2 -destination node1 -aggregate-list <aggregate_list> -ndo-controller-upgrade true -override-vetoes true -override-destination-checks true</pre>

Retire node2

To retire node2, you shut down node2 correctly and then remove it from the rack or chassis.

Resume the system controller replace operation

Steps

1. Resume the operation:

```
system controller replace resume
```

The node halts automatically.

Remove the AFF A800 or AFF C800 controller module


Remove the cable management device from the existing module and move the controller slightly out of the chassis.

Before you begin

If you are not already grounded, correctly ground yourself.

Steps

1. Prepare to remove the controller module:
 - a. On the front of the chassis, use your thumbs to firmly push each drive in to the top and bottom disk bays until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.



- b. Go to the rear of the chassis.
2. Unplug the node2 controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

5. Remove the cable management device from the controller module and set it aside.
6. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

After you finish

You can decommission node2 after the upgrade is completed. See [Decommission the old system](#).

Stage 5. Install the replacement system modules on node2

Install the replacement system modules on node2

Install AFF A30, AFF A50, AFF C30, or AFF C60 module on node2

Install the replacement controller module that you received for the upgrade on node2. Node2 is controller B located in the bottom half of the chassis when looking at the controllers from the rear of the system.

Steps

1. Position the storage system onto the rails in the middle of the cabinet or telco rack, and then support the storage system from the bottom and slide it into place.



Don't completely insert the controller module in the chassis until instructed to do so later in the procedure.

2. Cable the management and console ports to the node2 controller module.



Because the chassis is already powered ON, node2 starts BIOS initialization followed by AUTOBOOT as soon as it is fully seated. To interrupt the node2 boot, before completely inserting the controller module into the slot, it is recommended that you connect the serial console and management cables to the node2 controller module.

3. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated. The locking latch rises when the controller module is fully seated. Close the cam handle to the locked position.



To avoid damaging the connectors, don't use excessive force when sliding the controller module into the chassis.

4. Connect the serial console as soon as the module is seated and be ready to interrupt AUTOBOOT of node2.
5. After you interrupt AUTOBOOT, node2 stops at the LOADER prompt. If you do not interrupt AUTOBOOT on time and node2 starts booting, wait for the prompt and press Ctrl-C to go into the boot menu. After the node stops at the boot menu, use option 8 to reboot the node and interrupt AUTOBOOT during reboot.
6. At the LOADER> prompt of node2, set the default environment variables:

```
set-defaults
```

7. Save the default environment variables settings:

```
saveenv
```


8. For two-node switchless configurations, verify that you have an X60132A, 4-port 10/25 GbE card in slot1 on node2. The X60132A card is required for cluster interconnect during the upgrade.

Install the AFF A90 or AFF A70 module on node2

Install the AFF A90 or AFF A70 controller module that you received for the upgrade on node2. Node2 is controller B located on the lower bay of the chassis when looking at the controllers from the rear of the system.

Steps

1. On the front of the chassis, use your thumbs to firmly push each drive in to the top and bottom disk bays until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.




a. Go to the rear of the chassis.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.



Don't completely insert the controller module in the chassis until instructed to do so later in the procedure.



3. Cable the management and console ports to the node2 controller module.



Because the chassis is already powered ON, node2 starts BIOS initialization followed by AUTOBOOT as soon as it is fully seated. To interrupt the node2 boot, before completely inserting the controller module into the slot, it is recommended that you connect the serial console and management cables to the node2 controller module.

4. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated. The locking latch rises when the controller module is fully seated. Close the cam handle to the locked position.



To avoid damaging the connectors, don't use excessive force when sliding the controller module into the chassis.


5. Connect the serial console as soon as the module is seated and be ready to interrupt AUTOBOOT of node2.
6. After you interrupt AUTOBOOT, node2 stops at the LOADER prompt. If you do not interrupt AUTOBOOT on time and node2 starts booting, wait for the prompt and press Ctrl-C to go into the boot menu. After the node stops at the boot menu, use option 8 to reboot the node and interrupt AUTOBOOT during reboot.
7. At the LOADER> prompt of node2, set the default environment variables:

```
set-defaults
```

8. Save the default environment variables settings:

```
saveenv
```

9. On the front of the chassis, use your thumbs to firmly push each drive in to the top and bottom disk bays until you feel a positive stop. This ensures that the drives are firmly seated against the chassis midplane.



- a. Go to the rear of the chassis.

Install the ASA A150, AFF A150, or FAS2820 controller module on node2

Install the ASA A150, AFF A150 or FAS2820 controller module that you received for the upgrade on node2. Node2 is controller B located on the right side of the chassis when looking at the controllers from the rear of the system.

Before you begin

- If you are not already grounded, correctly ground yourself.
- Disconnect all the cables, including console, management, SAS storage, and data network cables, from the controller being removed.

Steps

1. Align the end of the controller module with bay B in the chassis, and then gently push the controller module halfway into the system.



Bay B is located on the right side of the chassis.



Don't completely insert the controller module in the chassis until you are instructed to do so later in the procedure.

2. Cable the management and console ports to the node2 controller module.



Because the chassis is already powered ON, node2 starts booting as soon as it is fully seated. To prevent node2 from booting, connect the console and management cables before fully inserting the controller module.

3. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latch rises when the controller module is fully seated.



To avoid damaging the connectors, don't use excessive force when sliding the controller module into the chassis.

4. Connect the serial console as soon as the module is seated and be ready to interrupt AUTOBOOT of node2.
5. After you interrupt AUTOBOOT, node2 stops at the LOADER prompt. If you do not interrupt AUTOBOOT on time and node2 starts booting, wait for the prompt and press Ctrl-C to go into the boot menu. After the node stops at the boot menu, use option 8 to reboot the node and interrupt AUTOBOOT during reboot.

Install the ASA A900, AFF A900, or FAS9500 NVRAM and controller modules on node2

Install the ASA A900, AFF A900, or FAS9500 NVRAM and controller modules that you received for the upgrade on node2. Node2 is controller B located on the right side of the chassis when looking at the controllers from the rear of the system.

You must note the following when performing the installation:

- Move all blank filler modules in slots 6-1 and 6-2 from the old NVRAM module to the new NVRAM module.
- Do NOT move the coredump device from the AFF A700 NVRAM module to the ASA A900 or AFF A900 NVRAM module.

- Move all flash cache modules installed in the FAS9000 NVRAM module to the FAS9500 NVRAM module.

Before you begin

If you are not already grounded, correctly ground yourself.

Install the ASA A900, AFF A900, or FAS9500 NVRAM module

Install the ASA A900, AFF A900, or FAS9500 NVRAM module in slot 6 of node2.

Steps

1. Align the NVRAM module with the edges of the chassis opening in slot 6.
2. Gently slide the NVRAM module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the NVRAM module in place.

Install the ASA A900, AFF A900, or FAS9500 controller module in node2

Install, cable, and connect the ASA A900, AFF A900, or FAS9500 controller module in node2.

Steps

1. Align the end of the controller module with bay B in the chassis, and then gently push the controller module halfway into the system.



The bay label is located on the chassis directly above the controller module.



Don't completely insert the controller module in the chassis until you are instructed to do so later in the procedure.

2. Cable the management and console ports to the node2 controller module.



Because the chassis is already powered ON, node2 starts booting as soon as it is fully seated. To avoid node2 booting, it is recommended that you connect the console and management cables to the node2 controller module before completely inserting the controller module into the slot.

3. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latch rises when the controller module is fully seated.



To avoid damaging the connectors, don't use excessive force when sliding the controller module into the chassis.

4. Connect the serial console as soon as the module is seated and be ready to interrupt AUTOBOOT of node2.
5. After you interrupt AUTOBOOT, node2 stops at the LOADER prompt. If you do not interrupt AUTOBOOT on time and node2 starts booting, wait for the prompt and press Ctrl-C to go into the boot menu. After the node stops at the boot menu, use option 8 to reboot the node and interrupt AUTOBOOT during reboot.
6. At the LOADER> prompt of node2, set the default environment variables:

```
set-defaults
```

7. Save the default environment variables settings:

```
saveenv
```

Netboot node2

After swapping the corresponding replacement node2 system modules, you might need to netboot them. The term netboot means that you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you put a copy of the ONTAP 9 boot image onto a web server that the system can access.

It isn't possible to check the version of ONTAP installed on the boot media of the replacement controller module unless it is installed in a chassis and powered ON. The ONTAP version on the replacement system boot media must be the same as the ONTAP version running on the old system that you are upgrading and both the primary and backup boot images must match. You can configure the images by performing a netboot followed by the `wipeconfig` command from the boot menu. If the controller module was previously used in another cluster, the `wipeconfig` command clears any residual configuration on the boot media.

You can also use the USB boot option to perform the netboot. See the Knowledge Base article [How to use the boot_recovery LOADER command for installing ONTAP for initial setup of a system](#).

Before you begin

- Verify that you can access a HTTP server with the system.
- Download the necessary system files for your system and the correct version of ONTAP from the *NetApp Support Site*. Refer to [References](#) to link to the *NetApp Support Site*.

About this task


You must netboot the new controllers if they do not have the same 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.

Steps

1. Refer to [References](#) to link to the *NetApp Support Site* to download the files used for performing the netboot of the system.
2. Download the appropriate ONTAP software from the software download section of the NetApp Support Site and store the `<ontap_version>_image.tgz` file on a web-accessible directory.
3. Change to the web-accessible directory and verify that the files you need are available.
4. Your directory listing should contain `<ontap_version>_image.tgz`.
5. Configure the netboot connection by choosing one of the following actions.



You must 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 using the following command at the boot environment prompt: <code>ifconfig e0M -auto</code>
Not running	<p>Manually configure the connection by using the following command at the boot environment prompt: <code>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></code></p> <p><i>filer_addr</i> is the IP address of the storage system (mandatory). <i>netmask</i> is the network mask of the storage system (mandatory). <i>gateway</i> is the gateway for the storage system (mandatory). <i>dns_addr</i> is the IP address of a name server on your network (optional). <i>dns_domain</i> is the Domain Name Service (DNS) domain name (optional).</p> <div>  <p>Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p> </div>

6. Perform netboot on node2:

```
netboot http://<web_server_ip/path_to_web_accessible_directory>/netboot/kernel
```



Do not interrupt the boot.

7. Wait for the node2 now running on the replacement controller module to boot and display the boot menu options as shown in the following output:

Please choose one of the following:

- (1) Normal Boot.
 - (2) Boot without /etc/rc.
 - (3) Change password.
 - (4) Clean configuration and initialize all disks.
 - (5) Maintenance mode boot.
 - (6) Update flash from backup config.
 - (7) Install new software first.
 - (8) Reboot node.
 - (9) Configure Advanced Drive Partitioning.
 - (10) Set Onboard Key Manager recovery secrets.
 - (11) Configure node for external key management.
- Selection (1-11)?

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

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

Disregard the following message:

This procedure is not supported for Non-Disruptive Upgrade on an HA pair

This note applies to nondisruptive ONTAP software upgrades, and not controller upgrades.



Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the incorrect image might install. This issue applies to all ONTAP releases. The netboot procedure combined with option (7) `Install new software` wipes the boot media and places the same ONTAP version on both image partitions.

9. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL:

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

The `<path_to_the_web-accessible_directory>` should lead to where you downloaded the `<ontap_version>_image.tgz` in [Step 2](#).

10. Complete the following substeps to reboot the controller module:

- 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}
```

- b. Enter `y` to reboot 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 must be restored.

11. Clear any previous configuration on the boot media.

- a. At the following prompt, run the `wipeconfig` command, and press the enter key:

Please choose one of the following:

- (1) Normal Boot.
 - (2) Boot without /etc/rc.
 - (3) Change password.
 - (4) Clean configuration and initialize all disks.
 - (5) Maintenance mode boot.
 - (6) Update flash from backup config.
 - (7) Install new software first.
 - (8) Reboot node.
 - (9) Configure Advanced Drive Partitioning.
 - (10) Set Onboard Key Manager recovery secrets.
 - (11) Configure node for external key management.
- Selection (1-11)? wipeconfig

b. When you see the message below, answer `yes`:

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

c. The node reboots to finish the `wipeconfig` and then stops at the boot menu.



Wait until the node stops at the boot menu after completing the `wipeconfig` operation.

12. Select maintenance mode 5 from the boot menu and enter `y` when you are prompted to continue with the boot.
13. Verify that the controller and chassis are configured as `ha`:

```
ha-config show
```

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

```
Chassis HA configuration: ha
Controller HA configuration: ha
```

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

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

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

15. Halt node2:

```
halt
```

Node2 should stop at the LOADER> prompt.

16. On node1, check the system date, time, and time zone:

```
date
```

17. On node2, check the date by using the following command at the boot environment prompt:

```
show date
```

18. If necessary, set the date on node2:

```
set date mm/dd/yyyy
```



Set the corresponding UTC date on node2.

19. On node2, check the time by using the following command at the boot environment prompt:

```
show time
```

20. If necessary, set the time on node2:

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



Set the corresponding UTC time on node2.

21. Set the partner system ID on node2:

```
setenv partner-sysid node1_sysid
```

For node2, the `partner-sysid` must be that of the node1 that you are upgrading.

a. Save the settings:

```
saveenv
```

22. On node2, at the LOADER prompt, verify the `partner-sysid` for node2:

```
printenv partner-sysid
```

Stage 6. Boot node2 with the replacement system modules

Cable node2 for shared cluster-HA and storage

If you are performing one of the following upgrades, you need to connect the cluster, HA, storage, data, and management connections that were previously connected to the node2 on the existing system to the newly installed node2 on the replacement system.

Existing system	Replacement system
AFF A250	AFF A30, AFF A50
AFF C250	AFF C30, AFF C60
AFF A800	AFF A70, AFF A90
AFF C800	AFF C80

Connect the e0M and BMC ports

If the existing system has a management port (e0M) and a BMC port, the e0M and BMC ports are combined and accessed through the "wrench" port on the replacement system. You must ensure that the e0M and BMC ports are connected to the same switch and subnet on the existing system before connecting to the replacement system.

If the...	Then...
e0M and BMC IP addresses are on the same IP subnet	Connect either the e0M or BMC port on the existing system to the "wrench" port on the replacement system.
e0M and BMC IP addresses are on different subnets	<ol style="list-style-type: none"> 1. Merge the e0M and BMC IP addresses into one IP subnet. 2. Connect either the e0M or BMC port on the existing system to the "wrench" port on the replacement system.

Connect to a two-node switchless cluster

The following tables show the switch port usage for two-node switchless cluster configurations.

Port type	AFF A800, AFF C800	AFF A90	AFF A70, AFF C80
Cluster	e0a	e1a	e1a
Cluster	e1a	e7a (Use e1b if there is no e7a)	e1b
HA	e0b	Don't connect	Don't connect
HA	e1b	Don't connect	Don't connect
SAS storage ports (if present and used)	Any available port	Any available port	Any available port
Ethernet storage ports for NS224 shelves	Any available port	Refer to Ethernet storage connectivity mapping	Refer to Ethernet storage connectivity mapping

Port type	AFF A250, AFF C250	AFF A30, AFF C60	AFF A50
Cluster	e0c	e1a (Use e1a for temporary cluster interconnect)	e1a (Use e1a for temporary cluster interconnect)

Port type	AFF A250, AFF C250	AFF A30, AFF C60	AFF A50
Cluster	e0d	e1b (Use e1b for temporary cluster interconnect)	e1b (Use e1b for temporary cluster interconnect)
HA	e0c HA port is shared with Cluster port	e4a on node1 is directly connected to e4a on node2 using a 100 GbE cable	e4a on node1 is directly connected to e4a on node2 using a 100 GbE cable
HA	e0d HA port is shared with Cluster port	e2a on node1 is directly connected to e2a on node2 using a 100 GbE cable If e2a isn't present or doesn't support 100 GbE, directly connect e4b on node1 to e4b on node2 using a 100 GbE cable.	e2a on node1 directly connected to e2a on node2 using a 100 GbE cable If e2a isn't present or doesn't support 100 GbE, directly connect e4b on node1 to e4b on node2 using a 100 GbE cable.
Ethernet storage port	Any available port	e3a, e3b	e3a, e3b
SAS storage port	Any available port	3a, 3b	3a, 3b

Connect to a switch-attached cluster

For a switch-attached cluster, check that you meet the following requirements for the AFF A30, AFF A50, AFF A70, AFF A90, AFF C30, AFF C60, or AFF C80 (replacement) node:

- The identical cluster ports on the replacement node are on the same switch. For example, on completion of the upgrade, e1a on node1 and e1a on node2 should be attached to one cluster switch. Similarly, the second cluster port from both nodes should be attached to the second cluster switch. Cross-connection of shared cluster-HA ports, where e1a from node1 is connected to switchA and e1a from node2 is connected to switchB, results in HA communication failures.
- The replacement node uses shared cluster-HA Ethernet ports.
- Verify that the cluster switches are installed with a reference configuration file (RCF) that supports shared cluster-HA ports:

1. Remove the existing configuration on the switch:

If your switch model is...	Go to...
Cisco Nexus	The Knowledge Base article How to clear configuration on a Cisco interconnect switch while retaining remote connectivity
Broadcom BES-53248	The Knowledge Base article How to clear configuration on a Broadcom interconnect switch while retaining remote connectivity

2. Configure and verify the switch setup:

If your switch model is...	Go to...
Cisco Nexus 9336C-FX2	Upgrade your Reference Configuration File (RCF)

If your switch model is...	Go to...
Broadcom BES-53248	Upgrade the Reference Configuration File (RCF)
NVIDIA SN2100	Install or upgrade the Reference Configuration File (RCF) script



If the cluster switch only supports 10/25 GbE speeds, you must use an X60130A, 4-port 10/25GbE card in slot1 or slot2 on the replacement system for cluster interconnect.

Boot node2 with the replacement system modules

Node2 with the replacement modules is now ready to boot. The supported replacement modules are listed in the [supported systems matrix](#).



You only move the console and management connections when you upgrade by swapping the system modules.

Steps

1. If you have NetApp Storage Encryption (NSE) drives installed, perform the following steps:



If you have not already done so earlier in the procedure, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use.

- a. Set `bootarg.storageencryption.support` to `true` or `false`:

If the following drives are in use...	Then...
NSE drives that conform to FIPS 140-2 Level 2 self-encryption requirements	<code>setenv bootarg.storageencryption.support true</code>
NetApp non-FIPS SEDs	<code>setenv bootarg.storageencryption.support false</code>



You cannot mix FIPS drives with other types of drives on the same node or HA pair. You can mix SEDs with non-encrypting drives on the same node or HA pair.

- b. Go to the special boot menu and select option (10) `Set Onboard Key Manager recovery secrets`.

Enter the passphrase and the backup information that you recorded earlier procedure. See [Manage storage encryption using the Onboard Key Manager](#).

2. Boot the node into the boot menu:

```
boot_ontap menu
```

3. When the node stops at the boot menu, reassign the old node2 disks to the replacement node2 by running the following command on node2:

```
boot_after_controller_replacement
```

After a short delay, you are prompted to enter the name of the node that is being replaced. If there are shared disks (also called Advanced Disk Partitioning (ADP) or partitioned disks), you are prompted to enter the node name of the HA partner.

These prompts might get buried in the console messages. If you do not enter a node name or enter an incorrect name, you are prompted to enter the name again.

```
If [localhost:disk.encryptNoSupport:ALERT]: Detected FIPS-certified
encrypting drive and, or, [localhost:diskown.errorDuringIO:error]:
error 3 (disk failed) on disk errors occur, perform the following steps:
```



1. Halt the node at the LOADER prompt.
2. Check and reset the storage encryption bootargs mentioned in [Step 1](#).
3. At the LOADER prompt, boot up:

```
boot_ontap
```

You can use the following example as a reference:

Expand the console output example

```
LOADER-A> boot_ontap menu
.
.
<output truncated>
.
All rights reserved.
*****
*                                     *
* Press Ctrl-C for Boot Menu. *
*                                     *
*****
.
<output truncated>
.
Please choose one of the following:

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 22/7

(22/7)                                Print this secret List
(25/6)                                Force boot with multiple filesystem
disks missing.
(25/7)                                Boot w/ disk labels forced to clean.
(29/7)                                Bypass media errors.
(44/4a)                               Zero disks if needed and create new
flexible root volume.
(44/7)                                Assign all disks, Initialize all
disks as SPARE, write DDR labels
.
.
<output truncated>
.
.
(wipeconfig)                          Clean all configuration on boot
```

```

device
(boot_after_controller_replacement) Boot after controller upgrade
(boot_after_mcc_transition)          Boot after MCC transition
(9a)                                Unpartition all disks and remove
their ownership information.
(9b)                                Clean configuration and
initialize node with partitioned disks.
(9c)                                Clean configuration and
initialize node with whole disks.
(9d)                                Reboot the node.
(9e)                                Return to main boot menu.

```

The boot device has changed. System configuration information could be lost. Use option (6) to restore the system configuration, or option (4) to initialize all disks and setup a new system. Normal Boot is prohibited.

Please choose one of the following:

```

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? boot_after_controller_replacement

```

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

```

.
.
<output truncated>
.
.
Controller Replacement: Provide name of the node you would like to
replace:<nodename of the node being replaced>
Controller Replacement: Provide High Availability partner of node1:
<nodename of the partner of the node being replaced>

```



```

Changing sysid of node nodel disks.
Fetched sanown old_owner_sysid = 536940063 and calculated old sys id
= 536940063
Partner sysid = 4294967295, owner sysid = 536940063
.
.
<output truncated>
.
.
varfs_backup_restore: restore using /mroot/etc/varfs.tgz
varfs_backup_restore: attempting to restore /var/kmip to the boot
device
varfs_backup_restore: failed to restore /var/kmip to the boot device
varfs_backup_restore: attempting to restore env file to the boot
device
varfs_backup_restore: successfully restored env file to the boot
device wrote key file "/tmp/rndc.key"
varfs_backup_restore: timeout waiting for login
varfs_backup_restore: Rebooting to load the new varfs
Terminated
<node reboots>

System rebooting...

.
.
Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...
Rebooting to load the restored env file...
.
System rebooting...
.
.
.
<output truncated>
.
.
.
.
WARNING: System ID mismatch. This usually occurs when replacing a
boot device or NVRAM cards!
Override system ID? {y|n} y
.
.
.

```

.

Login:



The system IDs shown in the preceding example are example IDs. The actual system IDs of the nodes that you are upgrading will be different.

Between entering node names at the prompt and the login prompt, the node reboots a few times to restore the environment variables, update firmware on the cards in the system, and for other ONTAP updates.

Verify the node2 installation

You must verify the node2 installation with the replacement system modules. Because there is no change to physical ports, you are not required to map the physical ports from the old node2 to the replacement node2.

About this task

After you boot node1 with the replacement system module, you verify that it is installed correctly. You must wait for node2 to join quorum and then resume the controller replacement operation.

At this point in the procedure, the operation pauses while node2 joins quorum.

Steps

1. Verify that node2 has joined quorum:

```
cluster show -node node2 -fields health
```

The output of the `health` field should be `true`.

2. This step applies to the following upgrade configurations. For all other system upgrades, skip this step and go to [Step 3](#):
 - Two node switchless clusters
 - Switch attached AFF A250 or AFF C250 systems upgrading to an AFF A50, AFF A30, AFF C30, or AFF C60 system.

If node2 doesn't join quorum automatically:

- a. Check the IPspace of ports e1a and e1b:

```
network port show
```

- b. If the IPspace isn't "Cluster", change the IPspace to "Cluster" on e1a and e1b:

```
network port modify -node <node_name> -port <port> -ipspace Cluster
```

- c. Verify that the IPspace of ports e1a and e1b is "Cluster":

```
network port show
```

d. Migrate node2 cluster LIFs to e1a and e1b:

```
network interface migrate -vserver Cluster -lif <cluster_lif1> -destination  
-node <node2_name> -destination-port <port_name>
```

3. Verify that node2 and node1 are part of the same cluster and that the cluster is healthy:

```
cluster show
```

4. Switch to advanced privilege mode:

```
set advanced
```

5. This step only applies to two-node switchless configuration upgrades from an AFF A250 or AFF C250 to an AFF A50, AFF A30, AFF C60, or AFF C30. For all other system upgrades, skip this step and go to [Step 6](#):

Verify that e4a, e2a, e1a, e1b ports or e4a, e4b, e1a, e1b ports are the cluster ports in "Cluster" broadcast domain.

The AFF A50, AFF A30, AFF C30, and AFF C60 systems share cluster and HA ports. You can safely migrate all cluster LIFs to e4a, e4b or e4a, e2a on node1 and node2:

a. List the home ports and current ports for all cluster LIFs:

```
network interface show -role Cluster -fields home-port,curr-port
```

b. On node1 and node2, migrate the cluster LIFs that are using e1a as the home port to e4a:

```
network interface migrate -vserver Cluster -lif <cluster_lif1> -destination  
-node <node> -destination-port e4a
```

c. On node1 and node2, modify the cluster LIFs migrated in [substep b](#) to use e4a as the home port:

```
network interface modify -vserver Cluster -lif <cluster_lif> -home-port e4a
```

d. Verify that the cluster is in quorum:

```
cluster show
```

e. Repeat [substep b](#) and [substep c](#) to migrate and modify the second cluster LIF on each node to e2a or e4b:

If e2a is present and is a 100GbE network port, this is the default second cluster port. If e2a isn't a 100GbE network port, ONTAP uses e4b as the second cluster and HA port.

f. Remove e1a and e1b from "Cluster" broadcast domain:

```
broadcast-domain remove-ports -broadcast-domain Cluster -ip-space Cluster  
-ports <node_name>:e1a
```

g. Verify that only cluster ports e4a, e2a or e4a, e4b are in "Cluster" broadcast domain

```
broadcast domain show
```

- h. Remove the cable connections between e1a node1 and e1a node2, and e1b node1 and e1b node2 to ensure only valid cluster-HA connections are used and there is no redundant connectivity.
6. Check the status of the controller replacement operation and verify that it is in a paused state and in the same state that it was in before node2 was halted to perform the physical tasks of installing new controllers and moving cables:

```
system controller replace show
```

```
system controller replace show-details
```

7. Resume the controller replacement operation:

```
system controller replace resume
```

8. The controller replacement operation pauses for intervention with the following message:

```
Cluster::*> system controller replace show
Node           Status           Error-Action
-----
Node2          Paused-for-intervention    Follow the instructions given
in
Step Details
Node1          None

Step Details:
-----
To complete the Network Reachability task, the ONTAP network
configuration must be manually adjusted to match the new physical
network configuration of the hardware. This includes:

1. Re-create the interface group, if needed, before restoring VLANs. For
detailed commands and instructions, refer to the "Re-creating VLANs,
ifgrps, and broadcast domains" section of the upgrade controller
hardware guide for the ONTAP version running on the new controllers.
2. Run the command "cluster controller-replacement network displaced-
vlans show" to check if any VLAN is displaced.
3. If any VLAN is displaced, run the command "cluster controller-
replacement network displaced-vlans restore" to restore the VLAN on the
desired port.
2 entries were displayed.
```



In this procedure, section *Re-creating VLANs, ifgrps, and broadcast domains* has been renamed *Restore network configuration on node2*.

9. With the controller replacement in a paused state, proceed to [Restore network configuration on node2](#).

Restore network configuration on node2

After you confirm that node2 is in quorum and can communicate with node1, verify that node1's VLANs, interface groups, and broadcast domains are seen on node2. Also, verify that all node2 network ports are configured in their correct broadcast domains.

About this task

For more information on creating and re-creating VLANs, interface groups, and broadcast domains, refer to [References](#) to link to the *Network Management* content.

Steps

1. List all the physical ports that are on upgraded node2:

```
network port show -node node2
```

All physical network ports, VLAN ports, and interface group ports on the node are displayed. From this output, you can see any physical ports that have been moved into the `Cluster` broadcast domain by ONTAP. You can use this output to aid in deciding which ports should be used as interface group member ports, VLAN base ports, or standalone physical ports for hosting LIFs.

2. List the broadcast domains on the cluster:

```
network port broadcast-domain show
```

3. List network port reachability of all ports on node2:

```
network port reachability show -node node2
```

You should see output similar to the following example. The port and broadcast names vary.

```
Cluster::> reachability show -node node1
(network port reachability show)
Node      Port      Expected Reachability      Reachability
Status
-----
Node1
    a0a      Default:Default      ok
    a0a-822   Default:822          ok
    a0a-823   Default:823          ok
    e0M       Default:Mgmt         ok
    e1a       Cluster:Cluster      ok
    e1b       -                    no-reachability
    e2a       -                    no-reachability
    e2b       -                    no-reachability
    e3a       -                    no-reachability
    e3b       -                    no-reachability
    e7a       Cluster:Cluster      ok
    e7b       -                    no-reachability
    e9a       Default:Default      ok
    e9a-822   Default:822          ok
    e9a-823   Default:823          ok
    e9b       Default:Default      ok
    e9b-822   Default:822          ok
    e9b-823   Default:823          ok
    e9c       Default:Default      ok
    e9d       Default:Default      ok
20 entries were displayed.
```

In the preceding example, node2 has booted and joined quorum after controller replacement. It has several ports that have no reachability and are pending a reachability scan.

4. Repair the reachability for each of the ports on node2 with a reachability status other than `ok` by using the following command, in the following order:

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

- a. Physical ports
- b. VLAN ports

You should see output like the following example:

```
Cluster ::> reachability repair -node node2 -port e9d
```

```
Warning: Repairing port "node2:e9d" may cause it to move into a
different broadcast domain, which can cause LIFs to be re-homed away
from the port. Are you sure you want to continue? {y|n}:
```

A warning message, as shown in the preceding example, is expected for ports with a reachability status that might be different from the reachability status of the broadcast domain where it is currently located. Review the connectivity of the port and answer `y` or `n` as appropriate.

Verify that all physical ports have their expected reachability:

```
network port reachability show
```

As the reachability repair is performed, ONTAP attempts to place the ports in the correct broadcast domains. However, if a port's reachability cannot be determined and does not belong to any of the existing broadcast domains, ONTAP will create new broadcast domains for these ports.

5. Verify port reachability:

```
network port reachability show
```

When all ports are correctly configured and added to the correct broadcast domains, the `network port reachability show` command should report the reachability status as `ok` for all connected ports, and the status as `no-reachability` for ports with no physical connectivity. If any port reports a status other than these two, perform the reachability repair and add or remove ports from their broadcast domains as instructed in [Step 4](#).

6. Verify that all ports have been placed into broadcast domains:

```
network port show
```

7. Verify that all ports in the broadcast domains have the correct maximum transmission unit (MTU) configured:

```
network port broadcast-domain show
```

8. Restore LIF home ports, specifying the Vserver and LIF home ports, if any, that need to be restored by using the following steps:

- a. List any LIFs that are displaced:

```
displaced-interface show
```

- b. Restore LIF home nodes and home ports:

```
displaced-interface restore-home-node -node node_name -vserver vserver_name
-lif-name LIF_name
```

9. Verify that all LIFs have a home port and are administratively up:

```
network interface show -fields home-port,status-admin
```

Restore key-manager configuration on node2

If you are using NetApp Aggregate Encryption (NAE) or NetApp Volume Encryption (NVE) to encrypt volumes on the system that you are upgrading, the encryption configuration must be synchronized to the new nodes. If you do not resynchronize the key-manager, when you relocate the node2 aggregates from the upgraded node1 to the upgraded node2 by using ARL, failures might occur because node2 does not have the required encryption keys to bring encrypted volumes and aggregates online.

About this task

Synchronize the encryption configuration to the new nodes by performing the following steps:

Steps

- 1. Run the following command from node2:

```
security key-manager onboard sync
```

- 2. Verify that the SVM-KEK key is restored to "true" on node2 before you relocate the data aggregates:

```
::> security key-manager key query -node node2 -fields restored -key
-type SVM-KEK
```

Example

```
::> security key-manager key query -node node2 -fields restored -key
-type SVM-KEK
```

node	vserver	key-server	key-id
restored			
-----	-----	-----	-----
node2	svm1	""	0000000000000000020000000000a008a81976
true			2190178f9350e071fbb90f000000000000000

Verify the RCF configuration on cluster switches

At this stage in the upgrade procedure, all data aggregates should be on node1. If you’re upgrading a configuration with switch-attached clusters, you need to validate that the cluster switch reference configuration file (RCF) supports the shared cluster/HA ports for the new controllers.

If you’re upgrading to a two-node switchless cluster configuration, you can skip this section and go to [Move non-root aggregates and NAS data LIFs back to node2](#).

Steps

1. Switch to advanced privilege mode:

```
set advanced
```

2. Check the status of "IC RDMA":

```
ha interconnect status show
```

In the output, the "IC RDMA Connection" should have the status Up.

If the "IC RDMA Connection" status is ...	Then...
Up	Go to Move non-root aggregates and NAS data LIFs back to node2.
Down	Go to Step 3.

3. Check the cluster ports and switch RCF.

For more information, see [Connect to a switch-attached cluster.](#)

4. Verify that the "IC RDMA Connection" status has changed to Up:

```
ha interconnect status show
```

What's next

[Move non-root aggregates and NAS data LIFs back to node2](#)

Move non-root aggregates and NAS data LIFs back to node2

After verifying the network configuration on node2, you need to relocate the NAS data LIFs owned by node2 from node1 to node2 and confirm that the SAN LIFs exist on node2.

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 the correct ports after you bring node2 online.

Steps

1. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

- Cluster quorum check
- System ID check

- Image version check
- Target platform check
- Network reachability check

The system pauses the operation at this stage in the network reachability check

2. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following checks:

- Cluster health check
- Cluster LIF status check

After performing these checks, the system relocates the non-root aggregates and NAS data LIFs back to node2, which is now running on the replacement controller.

The controller replacement operation pauses after the resource relocation is complete.

3. Check the status of the aggregate relocation and NAS data LIF move operations:

```
system controller replace show-details
```

If the controller replacement procedure is paused, check and correct the error, if any, and then issue `resume` to continue the operation.

4. If necessary, restore and revert displaced LIFs or manually migrate and modify the node2 LIFs that failed to relocate automatically to node2.

Restore and revert displaced LIFs

- a. List any displaced LIFs:

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

- b. If any LIFs are displaced, restore the home node back to node2:

```
cluster controller-replacement network displaced-interface  
restore-home-node -node <node2_nodename> -vserver <vserver name>  
-lif-name <lif_name>
```

Manually migrate and modify LIFs

- a. Migrate the LIFs that failed to relocate automatically to node2:

```
network interface migrate -vserver <vserver name> -lif <lif_name>  
-destination-node <node2_nodename> -destination-port  
<port_on_node2>
```

- b. Modify the home node and home port for the migrated LIFs:

```
network interface modify -vserver <vserver_name> -lif  
<data_lif_name> -home-node <node2_nodename> -home-port  
<home_port>
```

5. Resume the operation to prompt the system to perform the required post-checks:

```
system controller replace resume
```

The system performs the following post-checks:

- Cluster quorum check
- Cluster health check
- Aggregates reconstruction check
- Aggregate status check
- Disk status check
- Cluster LIF status check
- Volume check

Stage 7. Complete the upgrade

Manage authentication using KMIP servers

Beginning with ONTAP 9.10.1, you can use Key Management Interoperability Protocol (KMIP) servers to manage authentication keys.

Steps

1. Add a new controller:

```
security key-manager external enable
```

2. Add the key manager:

```
security key-manager external add-servers -key-servers  
key_management_server_ip_address
```

3. Verify that the key management servers are configured and available to all nodes in the cluster:

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

4. Restore the authentication keys from all linked key management servers to the new node:

```
security key-manager external restore -node new_controller_name
```

Confirm that the new controllers are set up correctly

To confirm the correct setup, you verify that the HA pair is enabled. You also verify that node1 and node2 can access each other's storage and that neither owns data LIFs belonging to other nodes on the cluster. In addition, you verify that all data aggregates are on their correct home nodes, and that the volumes for both nodes are online. If one of the new nodes has a unified target adapter, you must restore any port configurations and you might need to change the use of the adapter.

Steps

1. After the post-checks of node2, the storage failover and cluster HA pair for the node2 cluster are enabled. When the operation is done, both nodes show as completed and the system performs some cleanup operations.
2. Verify that storage failover is enabled:

```
storage failover show
```

The following example shows the output of the command when storage failover is enabled:

```
cluster::> storage failover show
```

Node	Partner	Takeover Possible	State Description
-----	-----	-----	-----
node1	node2	true	Connected to node2
node2	node1	true	Connected to node1

3. Verify that node1 and node2 belong to the same cluster by using the following command and examining the output:

```
cluster show
```

4. Verify that node1 and node2 can access each other's storage by using the following command and examining the output:

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

5. Verify that neither node1 nor node2 owns data LIFs home-owned by other nodes in the cluster by using the following command and examining the output:

```
network interface show
```

If neither node1 or node2 owns data LIFs home-owned by other nodes in the cluster, revert the data LIFs to their home owner:

```
network interface revert
```

6. Verify that the aggregates are owned by their respective home nodes.

```
storage aggregate show -owner-name <node1>
```

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

7. Determine whether any volumes are offline:

```
volume show -node <node1> -state offline
```

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

8. If any volumes are offline, compare them with the list of offline volumes that you captured in the section [Prepare the nodes for upgrade](#), and bring online any of the offline volumes, as required, by using the following command, once for each volume:

```
volume online -vserver <vserver_name> -volume <volume_name>
```

9. Install new licenses for the new nodes by using the following command 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 multiple licenses at once, separating each license key by a comma.

10. Remove all of the old licenses from the original nodes by using one of the following commands:

```
system license clean-up -unused -expired
```

```
system license delete -serial-number <node_serial_number> -package  
<licensable_package>
```

- Delete all expired licenses:

```
system license clean-up -expired
```

- Delete all unused licenses:

```
system license clean-up -unused
```

- Delete a specific license from a cluster by using 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.

11. Verify that the licenses are correctly installed by using the following command and examining its output:

```
system license show
```

You can compare the output with the output that you captured in the [Prepare the nodes for upgrade](#) section.

12. If self-encrypting drives are being used in the configuration and you have set the `kmip.init.maxwait` variable to `off` (for example, in *Boot node2 with the replacement system modules*, [Step 1](#)), you must unset the variable:

```
set diag; systemshell -node <node_name> -command sudo kenv -u -p  
kmip.init.maxwait
```

13. Configure the SPs by using the following command on both nodes:

```
system service-processor network modify -node <node_name>
```

Refer to [References](#) to link to the *System Administration Reference* for information about the SPs and the *ONTAP 9 Commands: Manual Page Reference* for detailed information about the `system service-processor network modify` command.

14. If you want to set up a switchless cluster on the new nodes, refer to [References](#) to link to the *NetApp*

Support Site and follow the instructions in *Transitioning to a two-node switchless cluster*.

After you finish

If Storage Encryption is enabled on node1 and node2, complete the section [Set up Storage Encryption on the new controller module](#). Otherwise, complete the section [Decommission the old system](#).

Set 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. You must 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 external show-status
```

```
security key-manager onboard show-backup
```

2. Add the key management servers listed in the previous step to the key management server list in the new controller.
 - a. Add the key management server:

```
security key-manager external add-servers -key-servers  
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 external 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 external enable
```

- 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 external restore -node new_controller_name
```

Set up NetApp Volume or Aggregate Encryption on the new controller module

If the replaced controller or high availability (HA) partner of the new controller uses NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE), you must configure the new controller module for NVE or NAE.

About this task

This procedure includes steps that are performed on the new controller module. You must enter the command on the correct node.

Onboard Key Manager

Configure NVE or NAE using the Onboard Key Manager.

Steps

1. Restore authentication keys from all linked key management servers to the new node:

```
security key-manager onboard sync
```

External Key Management

Configure NVE or NAE using External Key Management.

Steps

1. Verify that the key management servers are still available, their status, and their authentication key information:

```
security key-manager key query -node node
```

2. Add the key management servers listed in the previous step to the key management server list in the new controller:

- a. Add the key management server:

```
security key-manager external add-servers -key-servers  
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 external 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 external enable
```

- 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 external restore
```

This command needs the OKM passphrase

For more information, see the Knowledge Base article [How to restore external key manager server configuration from the ONTAP boot menu](#).

After you finish

Check if any volumes were taken offline because authentication keys were not available or EKM servers could not be reached. Bring those volumes back online by using the `volume online` command.

After you finish

Check if any volumes were taken offline because authentication keys were not available or External Key Management servers could not be reached. Bring those volumes back online by using the `volume online` command.

Decommission 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. Refer to [References](#) to link to the *NetApp Support Site* and log in.
2. Select **Products > My Products** from the menu.
3. On the **View Installed Systems** page, choose which **Selection Criteria** you want to use to display information about your system.

You can choose one of the following to locate your system:

- Serial Number (located on the back of the unit)
- Serial Numbers for My Location

4. Select **Go!**

A table displays cluster information, including the serial numbers.

5. Locate the cluster in the table and select **Decommission this system** from the Product Tool Set drop-down menu.

Resume SnapMirror operations

You can resume SnapMirror transfers that were quiesced before upgrade and resume the SnapMirror relationships. The updates are on schedule after the upgrade is completed.

Steps

1. Verify the SnapMirror status on the destination:

```
snapmirror show
```

2. Resume the SnapMirror relationship:

```
snapmirror resume -destination-vserver vserver_name
```

Troubleshoot

Aggregate relocation failures

Aggregate relocation (ARL) might fail at different points during the upgrade.

Check for aggregate relocation failure

During the procedure, ARL might fail in Stage 2, Stage 3, or Stage 5.

Steps

1. Enter the following command and examine the output:

```
storage aggregate relocation show
```

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, refer to [References](#) to link to the *ONTAP 9 Commands: Manual Page Reference*.

Aggregates originally on node1 are owned by node2 after completion of the upgrade

At the end of the upgrade procedure, node1 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 correctly, that is, they have node2 as their home node instead of node1, under the following circumstances:

- During Stage 3, when aggregates are relocated from node2 to node1.

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 is left behind on node2.

- After Stage 4, when node2 is replaced with the new system modules.

When node2 is replaced, `aggr_node_1` will come online with node1 as its home node instead of node2.

You can fix the incorrect ownership problem after Stage 6, after you have enabled storage failover by completing the following steps:

Steps

1. Get a list of aggregates:

```
storage aggregate show -nodes node2 -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 [Prepare the nodes for upgrade](#) and compare it with the output of the above command.

2. Compare the output of Step 1 with the output you captured for node1 in the section [Prepare the nodes for upgrade](#) and note any aggregates that were not correctly relocated.
3. Relocate the aggregates left behind on node2:

```
storage aggregate relocation start -node node2 -aggr aggr_node_1 -destination node1
```

Do not use the `-ndo-controller-upgrade` parameter during this relocation.

4. Verify that node1 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 node1 as home owner can be relocated to node1 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 the pre-check phase

Node1 or node2 crashes before the pre-check phase with HA pair still enabled

If either node1 or node2 crashes before the pre-check phase, no aggregates have been relocated yet and the HA pair 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.

Reboots, panics, or power cycles during first resource-release phase

Node1 crashes during the first resource-release phase with HA pair still enabled

Some or all aggregates have been relocated from node1 to node2, and HA pair is still enabled. Node2 takes over node1's root volume and any non-root aggregates that were not relocated.

About this task

Ownership of aggregates that were relocated look the same as the ownership of non-root aggregates that were taken over because the home owner has not changed.

When node1 enters the `waiting for giveback` state, node2 gives back all of the node1 non- root aggregates.

Steps

1. After node1 is booted up, all the non-root aggregates of node1 have moved back to node1. You must perform a manual aggregate relocation of the aggregates from node1 to node2:

```
storage aggregate relocation start -node node1 -destination node2 -aggregate  
-list * -ndocontroller-upgrade true
```
2. Continue with the node-pair upgrade procedure.

Node1 crashes during the first resource-release phase while HA pair is disabled

Node2 does 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.

Node2 fails during the first resource-release phase with HA pair still enabled

Node1 has relocated some or all of its aggregates to node2. The HA pair is enabled.

About this task

Node1 takes over all of node2's aggregates as well as any of its own aggregates that it had relocated to node2. When node2 boots up, the aggregate relocation is completed automatically.

Steps

1. Bring up node2.
2. Continue with the node-pair upgrade procedure.

Node2 crashes during the first resource-release phase and after HA pair is disabled

Node1 does not take over.

Steps

1. Bring up node2.

A client outage occurs for all aggregates while node2 is booting up.
2. Continue the with rest of the node-pair upgrade procedure.

Reboots, panics, or power cycles during the first verification phase

Node2 crashes during the first verification phase with HA pair disabled

Node1 does not take over following a node2 crash as the HA pair is already disabled.

Steps

1. Bring up node2.

A client outage occurs for all aggregates while node2 is booting up.

2. Continue with the node-pair upgrade procedure.

Node1 crashes during the first verification phase with HA pair disabled

Node2 does 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.

Reboots, panics, or power cycles during first resource-regain phase

Node2 crashes during the first resource-regain phase during aggregate relocation

Node2 has relocated some or all of its aggregates from node1 to node1. Node1 serves data from aggregates that were relocated. The HA pair is disabled and hence there is no takeover.

About this task

There is client outage for aggregates that were not relocated. On booting up node2, the aggregates of node1 are relocated to node1.

Steps

1. Bring up node2.
2. Continue with the node-pair upgrade procedure.

Node1 crashes during the first resource-regain phase during aggregate relocation

If node1 crashes while node2 is relocating aggregates to node1, the task continues after node1 boots up.

About this task

Node2 continues to serve remaining aggregates, but aggregates that were already relocated to node1 encounter client outage while node1 is booting up.

Steps

1. Bring up node1.
2. Continue with the controller upgrade.

Reboots, panics, or power cycles during post-check phase

Node1 or node2 crashes during the post-check phase

The HA pair is disabled hence this is no takeover. There is a client outage for aggregates belonging to the node that rebooted.

Steps

1. Bring up the node.
2. Continue with the node-pair upgrade procedure.

Reboots, panics, or power cycles during second resource-release phase

Node1 crashes during the second resource-release phase

If node1 crashes while node2 is relocating aggregates, the task continues after node1 boots up.

About this task

Node2 continues to serve remaining aggregates but aggregates that were already relocated to node1 and node1's own aggregates encounter client outages while node1 is booting.

Steps

1. Bring up node1.
2. Continue with the controller upgrade procedure.

Node2 crashes during the second resource-release phase

If node2 crashes during aggregate relocation, node2 is not taken over.

About this task

Node1 continues to serve the aggregates that have been relocated, but the aggregates owned by node2 encounter client outages.

Steps

1. Bring up node2.
2. Continue with the controller upgrade procedure.

Reboots, panics, or power cycles during the second verification phase

Node1 crashes during the second verification phase

If node1 crashes during this phase, takeover does not happen because the HA pair is already disabled.

About this task

There is a client outage for all aggregates until node1 reboots.

Steps

1. Bring up node1.
2. Continue with the node-pair upgrade procedure.

Node2 crashes during the second verification phase

If node2 crashes during this phase, takeover does not happen. Node1 serves data from the aggregates.

About this task

There is an outage for non-root aggregates that were already relocated until node2 reboots.

Steps

1. Bring up node2.
2. Continue with the node-pair upgrade procedure.

Issues that can arise in multiple stages of the procedure

Some issues can occur during 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	
		Possible	State Description
node1	node2	false	Unknown
node2	node1	false	Node owns partner aggregates as part of the non-disruptive head upgrade procedure. Takeover is not possible: Storage failover is disabled.

The following example shows the output of the `storage failover show` command after a reboot or panic:

```
cluster::> storage failover show
```

Node	Partner	Takeover	
		Possible	State Description
node1	node2	-	Unknown
node2	node1	false	Waiting for node1, Partial giveback, Takeover is not possible: Storage failover is disabled

Although the output says that a node is in partial giveback and that storage failover is disabled, you can disregard this message.

Steps

No action is required; continue with the node-pair upgrade procedure.

LIF migration failure

After you migrate LIFs, they might not come online after migration in Stage 2, Stage 3, or Stage 5.

Steps

1. Verify that the 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.

2. Check the physical connectivity of the network cable if the physical state of the port is `down`.

References

When performing the procedures in this content, you might need to consult reference content or go to reference websites.

Reference content

Content specific to this upgrade are listed in the table below.

Content	Description
Administration overview with the CLI	Describes how to administer ONTAP systems, shows you how to use the CLI interface, how to access the cluster, how to manage nodes, and much more.
Decide whether to use System Manager or the ONTAP CLI for cluster setup	Describes how to set up and configure ONTAP.
Disk and aggregate management with the CLI	Describes how to manage ONTAP physical storage using the CLI. It shows you how to create, expand, and manage aggregates, how to work with Flash Pool aggregates, how to manage disks, and how to manage RAID policies.
HA pair management	Describes how to install and manage high-availability clustered configurations, including storage failover and takeover/giveback.
Logical storage management with the CLI	Describes how to efficiently manage your logical storage resources, using volumes, FlexClone volumes, files, and LUNs, FlexCache volumes, deduplication, compression, qtrees, and quotas.
MetroCluster Management and Disaster Recovery	Describes how to perform MetroCluster switchover and switchback operations, both in planned maintenance operations, or in the event of a disaster.

Content	Description
MetroCluster Upgrade and Expansion	Provides procedures for upgrading controller and storage models in the MetroCluster configuration, transitioning from a MetroCluster FC to a MetroCluster IP configuration, and expanding the MetroCluster configuration by adding additional nodes.
Network Management	Describes how to configure and manage physical and virtual network ports (VLANs and interface groups), LIFs, routing, and host-resolution services in clusters; optimize network traffic by load balancing; and monitor the cluster by using SNMP.
ONTAP 9.0 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.0 commands.
ONTAP 9.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.1 commands.
ONTAP 9.2 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.2 commands.
ONTAP 9.3 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.3 commands.
ONTAP 9.4 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.4 commands.
ONTAP 9.5 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.5 commands.
ONTAP 9.6 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.6 commands.
ONTAP 9.7 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.7 commands.
ONTAP 9.8 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.8 commands.
ONTAP 9.9.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.9.1 commands.
ONTAP 9.10.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.10.1 commands.
SAN management with the CLI	Describes how to configure and manage LUNs, igroups, and targets using the iSCSI and FC protocols, and namespaces and subsystems using the NVMe/FC protocol.
SAN configuration reference	Contains information about FC and iSCSI topologies and wiring schemes.
Upgrade by moving volumes or storage	Describes how to quickly upgrade controller hardware in a cluster by moving storage or volumes. Also describes how to convert a supported model to a disk shelf.
Upgrade ONTAP	Contains instructions for downloading and upgrading ONTAP.

Content	Description
Use "system controller replace" commands to upgrade controller hardware introduced in ONTAP 9.15.1 and later	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers introduced in ONTAP 9.15.1 and later by using "system controller replace" commands.
Use "system controller replace" commands to upgrade controller models in the same chassis	Describes the aggregate relocation procedures needed to non-disruptively upgrade a system, keeping the old system chassis and disks.
Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers running ONTAP 9.8 by using "system controller replace" commands.
Use aggregate relocation to manually upgrade controller hardware running ONTAP 9.8 or later	Describes the aggregate relocation procedures needed to perform manual non-disruptive controller upgrades running ONTAP 9.8 or later.
Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.5 to ONTAP 9.7	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers running ONTAP 9.5 to ONTAP 9.7 by using "system controller replace" commands.
Use aggregate relocation to manually upgrade controller hardware running ONTAP 9.7 or earlier	Describes the aggregate relocation procedures needed to perform manual non-disruptive controller upgrades running ONTAP 9.7 or earlier.

Reference sites

The [NetApp Support Site](#) also contains documentation about network interface cards (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.

Access [ONTAP 9 documentation](#).

Access the [Active IQ Config Advisor](#) tool.

Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later

Learn about this ARL upgrade procedure

There are several aggregate relocation (ARL) methods for upgrading controller hardware. This procedure describes how to upgrade the controller hardware using aggregate relocation (ARL) with "system controller replace commands" on systems running ONTAP 9.8 or later.

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 confirm 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.

Terminology used in this information

In this information, the original nodes are called "node1" and "node2", and the new nodes are called "node3"

and "node4". During the described procedure, node1 is replaced by node3, and node2 is replaced by node4.

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

Important information:

- This procedure is complex and assumes that you have advanced ONTAP administration skills. You also must read and understand the [guidelines for upgrading controllers with ARL](#) and the [ARL upgrade sequence](#) before beginning the upgrade.
- 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.
- You can use ARL to perform a non-disruptive simplified controller upgrade to a new controller running a later ONTAP version than the version running on the cluster you are upgrading. The ONTAP version combinations for old and new controllers are determined by the ONTAP software release NDU cadence model. For example, if you have a controller running ONTAP 9.8, and that is the last supported version for that controller, you can upgrade to a new controller running an ONTAP version later than ONTAP 9.8.

This upgrade procedure primarily applies to upgrade scenarios where the controller model you are replacing does not support later ONTAP versions and the new controller does not support earlier ONTAP versions.

- 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 high-availability (HA) pair in the cluster.
- This procedure applies to FAS systems and AFF systems.
- This procedure applies to systems running 4-node NetApp MetroCluster configuration or higher. Since MetroCluster configuration sites can be at two physically different locations, the automated controller upgrade must be carried out individually at each MetroCluster site for an HA pair.
- For non-MetroCluster systems, such as HA clusters, the ARL upgrade is the only supported procedure.
- If you are upgrading from an AFF A320 system, you can use volume moves to upgrade controller hardware or contact technical support. Refer to [References](#) to link to *Upgrade by moving volumes or storage*.

Automate the controller upgrade process

During a controller upgrade, the controller is replaced with another controller running a newer or more powerful platform. Earlier versions of this content contained instructions for a nondisruptive controller update process that was comprised of entirely manual steps. This content provides the steps for the new automated procedure which utilizes automatic network port reachability checks to further simplify the controller upgrade experience.

The manual process was lengthy and complex but in this simplified procedure you can implement a controller update using aggregate relocation, enabling more efficient nondisruptive upgrades for HA pairs. There are significantly fewer manual steps, especially around validation, collection of information, and post checks.

Decide whether to use this aggregate relocation procedure

There are several aggregate relocation (ARL) methods for upgrading controller hardware. This procedure describes how to upgrade the controller hardware using aggregate relocation (ARL) with "system controller replace commands" on systems running ONTAP 9.8 or later. You should only use this complex procedure if you're an experienced ONTAP administrator.

To help you decide if this ARL procedure is suitable for your controller hardware upgrade, you should review all of the following circumstances for supported upgrades:

- You're running ONTAP 9.8 or later.
- You don't want to add the new controllers as a new HA pair to the cluster and migrate the data using volume moves.
- You're experienced in administering ONTAP and are comfortable with the risks of working in the diagnostic privilege mode.
- If you're upgrading a MetroCluster configuration, it's a four-node or higher FC configuration, and all nodes are running ONTAP 9.8 or later.

For upgrading MetroCluster IP configurations, refer to [References](#) to link to *MetroCluster Upgrade and Expansion*.



- If you're upgrading a system by swapping controller modules in the same chassis, such as AFF A800 or AFF C800, NetApp strongly recommends using the upgrade procedure that [upgrades controller models using ARL, keeping the existing system chassis and disks](#). This ARL procedure includes the steps that ensure the internal disks remain secure in the chassis when you remove and install the controllers during the upgrade procedure.

[Learn about the supported system upgrade combinations using ARL, keeping the existing system chassis and disks.](#)

- You can use NetApp Storage Encryption (NSE), NetApp Volume Encryption (NVE), and NetApp Aggregate Encryption (NAE) with this procedure.

Supported system upgrade combinations

The following table shows the supported systems matrix for performing a controller upgrade using this ARL procedure.

Old controller	Replacement controller
FAS8020 ³ , FAS8040 ³ , FAS8060, FAS8080	FAS8200, FAS8300, FAS8700, FAS9000
FAS8060 ⁴ , FAS8080 ⁴	FAS9500
AFF8020 ³ , AFF8040 ³ , AFF8060, AFF8080	AFF A300, AFF A400, AFF A700, AFF A800 ¹
AFF8060 ⁴ , AFF8080 ⁴	AFF A900
FAS8200	FAS8300 ² , FAS8700, FAS9000, FAS9500
FAS8300, FAS8700, FAS9000	FAS9500

Old controller	Replacement controller
AFF A300	AFF A400 ² , AFF A700, AFF A800 ¹ , AFF A900
AFF A320 ⁴	AFF A400
AFF A400, AFF A700	AFF A900



If your controller upgrade model combination isn't in the above table, contact technical support.

¹For the additional steps needed for AFF A800 systems, go to the step that references A800 in section [Reassign node1 disks to node3, Step 9](#), or [Reassign node2 disks to node4, Step 9](#).

²If you're upgrading from an AFF A300 to an AFF A400 or an FAS8200 to an FAS8300 system in a two-node switchless cluster configuration, you must pick temporary cluster ports for the controller upgrade. The AFF A400 and FAS8300 systems come in two configurations, as an Ethernet bundle where the mezzanine card ports are Ethernet type and as an FC bundle where the mezzanine ports are FC type.

- For an AFF A400 or an FAS8300 with an Ethernet type configuration, you can use any of the two mezzanine ports as temporary cluster ports.
- For an AFF A400 or an FAS8300 with an FC type configuration, you must add a four-port 10GbE network interface card (part number X1147A) to provide temporary cluster ports.
- After you complete a controller upgrade by using temporary cluster ports, you can nondisruptively migrate cluster LIFs to e3a and e3b, 100GbE ports on an AFF A400 system, and e0c and e0d, 100GbE ports on an FAS8300 system.

³For FAS8020, FAS8040, AFF8020, and AFF8040 system upgrades to the target replacement controllers listed in the table above, the replacement controllers must be running same ONTAP version as the old controller. Note that FAS8020, FAS8040, AFF8020, and AFF8040 systems do not support ONTAP versions later than ONTAP 9.8.

⁴The following table shows the minimum and later supported ONTAP versions for these controller upgrade combinations.

Old controller		Replacement controller	
System	ONTAP version	System	ONTAP version
AFF A320	9.9.1 or later	AFF A400	9.9.1 or later
AFF8060	9.8P13 or later patches	AFF A900	9.10.1 to 9.12.1
AFF8080	9.8P10 or later patches	AFF A900	9.10.1 to 9.12.1
FAS8060	9.8P13 or later patches	FAS9500	9.10.1P3 to 9.12.1
FAS8080	9.8P12 or later patches	FAS9500	9.10.1P3 to 9.12.1

For the upgrade combinations shown in the preceding table:



- It isn't required to use the same ONTAP version on the existing and replacement controllers. The ONTAP software upgrade is performed with the controller upgrade.
- When upgrading, you must install a replacement controller with a supported ONTAP version and patch level.
- It isn't possible to cancel or back out of a controller upgrade after you start the procedure and upgrade the first node.

Choose a different hardware upgrade procedure

- [Review the alternative ARL methods available for upgrading controller hardware.](#)
- If you prefer a different method of upgrading the controller hardware and are willing to do volume moves, refer to [References](#) to link to *Upgrade by moving volumes or storage*.

Related information

Refer to [References](#) to link to the *ONTAP 9 Documentation*.

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 need the following tools to perform the upgrade:

- Grounding strap
- #2 Phillips screwdriver

Go to the [References](#) section to access the list of reference documents and reference sites required for this upgrade

Guidelines for upgrading controllers with ARL

To understand whether you can use ARL to upgrade a pair of controllers running ONTAP 9.8 or later depends on the platform and the configuration of both the original and replacement controllers.

Supported upgrades for ARL

When you upgrade a pair of nodes using this ARL procedure for ONTAP 9.8 or later, you must verify that ARL can be performed on the original and replacement controllers.

You must check the size of all defined aggregates and number of disks supported by the original system. You must then compare the aggregate sizes and number of disks supported to the aggregate size and number of disks supported by the new system. Refer to [References](#) to link to the *Hardware Universe* where this information is available. 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 must 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. For more information about cluster mixing rules, refer to

[References](#) to link to the *Hardware Universe*.



If you are upgrading a system that supports internal drives (for example, an FAS2700 or AFF A250) but does NOT have internal drives, refer to [References](#) and use the procedure in the *Aggregate Relocation to Manually Upgrade Controller Hardware* content that is correct for your version of ONTAP.

If you have a system with more than two cluster ports per node, such as an FAS8080 or an AFF8080 system, before you start the upgrade, you must migrate and re-home the cluster LIFs to two cluster ports per node. If you perform the controller upgrade with more than two cluster ports per node, cluster LIFs might be missing on the new controller after the upgrade.

Controller upgrade using ARL is supported on systems configured with SnapLock Enterprise and SnapLock Compliance volumes.

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.

Upgrades not supported for ARL

You cannot perform the following upgrades:

- To replacement controllers that do not support the disk shelves connected to the original controllers

Refer to [References](#) to link to the *Hardware Universe* for disk-support information.

- To entry level controllers with internal drives, for example: an FAS 2500.

If you want to upgrade entry level controllers with internal drives, refer to [References](#) to link to *Upgrade by moving volumes or storage* and go to the procedure *Upgrading a pair of nodes running clustered Data ONTAP by moving volumes*.

MetroCluster FC configuration

In a MetroCluster FC configuration, you must replace the disaster recovery/failover site nodes as soon as possible. Mismatch in controller models within a MetroCluster is not supported because controller model mismatch can cause disaster recovery mirroring to go offline. Use the `-skip-metrocluster-check true` command to bypass MetroCluster checks when you are replacing nodes at the second site.

Troubleshoot

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.

If any problems occur, refer to the [Troubleshoot](#) section at the end of the procedure for more information and possible solutions. Information about the failures that can occur is listed by the phase of the procedure in the [ARL upgrade sequence](#).

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

Verify the health of the MetroCluster configuration

Before starting an upgrade on a Fabric MetroCluster configuration, you must check the health of the MetroCluster configuration to verify correct operation.

Steps

1. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
metrocluster_siteA::*> metrocluster check run
```

The operation runs in the background.

2. After the `metrocluster check run` operation completes, view the results:

```
metrocluster check show
```

After approximately five minutes, the following results are displayed:

```
metrocluster_siteA::*> metrocluster check show
Last Checked On: 4/7/2019 21:15:05
Component          Result
-----
nodes               ok
lifs                ok
config-replication ok
aggregates          warning
clusters            ok
connections         not-applicable
volumes             ok
7 entries were displayed.
```

3. Check the status of the running MetroCluster check operation:

```
metrocluster operation history show -job-id 38
```

4. Verify that there are no health alerts:

```
system health alert show
```

Check for MetroCluster configuration errors

You can use the Active IQ Config Advisor tool available from the NetApp Support Site to check for common configuration errors.

If you do not have a MetroCluster configuration, you can skip this section.

About this task

Active IQ Config Advisor is a configuration validation and health check tool. You can deploy it at both secure sites and non-secure sites for data collection and system analysis.



Support for Config Advisor is limited, and available only online.

1. Download the [Active IQ Config Advisor](#) tool.
2. Run Active IQ Config Advisor, reviewing the output and following its recommendations to address any issues.

Verify switchover, healing, and switchback

You should verify the switchover, healing, and switchback operations of the MetroCluster configuration.

Refer to [References](#) to link to the *MetroCluster Management and Disaster Recovery* content and use the procedures mentioned for negotiated switchover, healing, and switchback.

Learn about the ARL upgrade sequence

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

Upgrade 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.

ARL upgrade sequence overview

During the procedure, you upgrade the original controller hardware with the replacement controller hardware, one controller at a time, taking advantage of the HA pair 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 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:

Stage	Description
Stage 1. Prepare for upgrade	<p>During Stage 1, you run prechecks and, if required, correct aggregate ownership. You must record certain information if you are managing storage encryption by using the OKM and you can choose to quiesce the SnapMirror relationships.</p> <p>Aggregate ownership at the end of Stage 1:</p> <ul style="list-style-type: none"> • Node1 is the home owner and current owner of the node1 aggregates. • Node2 is the home owner and current owner of the node2 aggregates.
Stage 2. Relocate and retire node1	<p>During Stage 2, you relocate node1 non-root aggregates and NAS data LIFs to node2. This process is largely automated; the operation pauses to enable you to check its status. You must manually resume the operation. If required, you relocate failed or vetoed aggregates. You record node1 information for use later in the procedure before retiring node1. You can also prepare to netboot node3 and node4 later in the procedure.</p> <p>Aggregate ownership at the end of Stage 2:</p> <ul style="list-style-type: none"> • Node2 is the current owner of node1 aggregates. • Node2 is the home owner and current owner of node2 aggregates.
Stage 3. Install and boot node3	<p>During Stage 3, you install and boot node3, check that the cluster and node-management ports from node1 come online on node3, reassign the node1 disks to node3, and verify the node3 installation. If you are using NetApp Volume Encryption (NVE), you restore key-manager configuration. If required, you set the FC or UTA/UTA2 configuration on node3. You also relocate the node1 NAS data LIFs and non-root aggregates from node2 to node3 and verify that the SAN LIFs exist on node3.</p> <p>Aggregate ownership at the end of Stage 3:</p> <ul style="list-style-type: none"> • Node3 is the home owner and current owner of node1 aggregates. • Node2 is the home owner and current owner of node2 aggregates.

Stage	Description
Stage 4. Relocate and retire node2	<p>During Stage 4, you relocate non-root aggregates and NAS data LIFs from node2 to node3. You also record node2 information for use later in the procedure before retiring node2.</p> <p>Aggregate ownership at the end of Stage 4:</p> <ul style="list-style-type: none"> • Node3 is the home owner and current owner of aggregates that originally belonged to node1. • Node2 is the home owner of node2 aggregates. • Node3 is the current owner of node2 aggregates.
Stage 5. Install and boot node4	<p>During Stage 5, you install and boot node4, check that the cluster and node-management ports from node2 come online on node4, reassign the node2 disks to node4, and verify the node4 installation. If you are using NVE, you restore key-manager configuration. If required, you set the FC or UTA/UTA2 configuration on node4. You also relocate node2 NAS data LIFs and non-root aggregates from node3 to node4 and verify that the SAN LIFs exist on node4.</p> <p>Aggregate ownership at the end of Stage 5:</p> <ul style="list-style-type: none"> • Node3 is the home owner and current owner of the aggregates that originally belonged to node1. • Node4 is the home owner and current owner of aggregates that originally belonged to node2.
Stage 6. Complete the upgrade	<p>During Stage 6, you confirm that the new nodes are set up correctly and, if the new nodes are encryption-enabled, you configure and set up Storage Encryption or NVE. You should also decommission the old nodes and resume the SnapMirror operations.</p>

Stage 1. Prepare for upgrade

Prepare the nodes for upgrade

The controller replacement process begins with a series of prechecks. You also gather information about the original nodes for use later in the procedure and, if required, determine the type of self-encrypting drives that are in use.

Steps

1. Begin the controller replacement process by entering the following command in the ONTAP command line:

```
system controller replace start -nodes node_names
```

- Beginning with ONTAP 9.10.1, the automated negotiated switchover (NSO) based upgrade procedure is the default for a four-node MetroCluster FC configuration. If you are upgrading a four-node MetroCluster FC configuration, when you issue the `system controller replace start` command, you must prevent the NSO based procedure initiating by setting the `-nso` parameter to `false`:



```
system controller replace start -nodes node_names -nso false
```

- You can only execute the `system controller replace start` command at the advanced privilege level:

```
set -privilege advanced
```

You will see the following output:

Warning:

1. Current ONTAP version is 9.x

Before starting controller replacement operation, ensure that the new controllers are running the version 9.x

2. Verify that NVMEM or NVRAM batteries of the new nodes are charged, and charge them if they are not. You need to physically check the new nodes to see if the NVMEM or NVRAM batteries are charged. You can check the battery status either by connecting to a serial console or using SSH, logging into the Service Processor (SP) or Baseboard Management Controller (BMC) 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.

3. If a controller was previously part of a different cluster, run `wipeconfig` before using it as the replacement controller.

Do you want to continue? {y|n}: y

2. Press `y`, you will see the following output:

Controller replacement operation: Prechecks in progress.

Controller replacement operation has been paused for user intervention.

The system runs the following prechecks; record the output of each precheck for use later in the procedure:

Precheck	Description
Cluster Health Check	Checks all the nodes in the cluster to confirm they are healthy.

Precheck	Description
MCC Cluster Check	<p>Checks if the system is a MetroCluster configuration. The operation automatically detects if it is a MetroCluster configuration or not and performs the specific prechecks and verification checks.</p> <p>Only 4-node MetroCluster FC configuration is supported. In the case of 2-node MetroCluster configuration and 4-node MetroCluster IP configuration, the check fails.</p> <p>If the MetroCluster configuration is in switched over state, the check fails.</p>
Aggregate Relocation Status Check	<p>Checks whether an aggregate relocation is already in progress. If another aggregate relocation is in progress, the check fails.</p>
Model Name Check	<p>Checks whether the controller models are supported for this procedure.</p> <p>If the models are not supported, the task fails.</p>
Cluster Quorum Check	<p>Checks that the nodes being replaced are in quorum. If the nodes are not in quorum, the task fails.</p>
Image Version Check	<p>Checks that the nodes being replaced run the same version of ONTAP.</p> <p>If the ONTAP image versions are different, the task fails.</p> <p>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. For instructions on how to upgrade ONTAP, refer to References to link to <i>Upgrade ONTAP</i>.</p>
HA Status Check	<p>Checks if both the nodes being replaced are in a high- availability (HA) pair configuration.</p> <p>If storage failover is not enabled for the controllers, the task fails.</p>
Aggregate Status Check	<p>If the nodes being replaced own aggregates for which they are not the home owner, the task fails.</p> <p>The nodes should not own any non-local aggregates.</p>
Disk Status Check	<p>If any nodes being replaced have missing or failed disks, the task fails.</p> <p>If any disks are missing, refer to References to link to <i>Disk and aggregate management with the CLI</i>, <i>Logical storage management with the CLI</i>, and <i>HA pair management</i> to configure storage for the HA pair.</p>
Data LIF Status Check	<p>Checks if any of the nodes being replaced have non- local data LIFs.</p> <p>The nodes should not contain any data LIFs for which they are not the home owner. If one of the nodes contains non-local data LIFs, the task fails.</p>
Cluster LIF Status	<p>Checks whether the cluster LIFs are up for both nodes. If the cluster LIFs are down, the task fails.</p>
ASUP Status Check	<p>If ASUP notifications are not configured, the task fails.</p> <p>You must enable ASUP before beginning the controller replacement procedure.</p>

Precheck	Description
CPU Utilization Check	Checks if the CPU utilization is more than 50% for any of the nodes being replaced. If the CPU usage is more than 50% for a considerable period of time, the task fails.
Aggregate Reconstruction Check	Checks if reconstruction is occurring on any data aggregates. If aggregate reconstruction is in progress, the task fails.
Node Affinity Job Check	Checks if any node affinity jobs are running. If node affinity jobs are running, the check fails.

- After the controller replacement operation is started and the prechecks are completed, the operation pauses enabling you to collect output information that you might need later when configuring node3.



If you have a system with more than two cluster ports per node, such as an FAS8080 or an AFF8080 system, before you start the upgrade, you must migrate and re-home the cluster LIFs to two cluster ports per node. If you perform the controller upgrade with more than two cluster ports per node, cluster LIFs might be missing on the new controller after the upgrade.

- Run the below set of commands as directed by the controller replacement procedure on the system console.

From the serial port connected to each node, run and save the output of the following commands individually:

```

° vservers services name-service dns show
° network interface show -curr-node local -role cluster,intercluster,node-
  mgmt,cluster-mgmt,data
° network port show -node local -type physical
° service-processor show -node local -instance
° network fcp adapter show -node local
° network port ifgrp show -node local
° system node show -instance -node local
° run -node local sysconfig
° storage aggregate show -node local
° volume show -node local
° storage array config show -switch switch_name
° system license show -owner local
° storage encryption disk show
° security key-manager onboard show-backup
° security key-manager external show
° security key-manager external show-status

```

◦ `network port reachability show -detail -node local`



If NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE) using the Onboard Key Manager (OKM) is in use, keep the key manager passphrase ready to complete the key manager resync later in the procedure.

5. If your system uses self-encrypting drives, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use on the HA pair that you are upgrading. ONTAP software supports two types of self-encrypting drives:

- FIPS-certified NetApp Storage Encryption (NSE) SAS or NVMe drives
- Non-FIPS self-encrypting NVMe drives (SED)



You cannot mix FIPS drives with other types of drives on the same node or HA pair.

You can mix SEDs with non-encrypting drives on the same node or HA pair.

[Learn more about supported self-encrypting drives.](#)

Correct aggregate ownership if an ARL precheck fails

If the Aggregate Status Check fails, you must return aggregates owned by the partner node to the home owner node and initiate the precheck process again.

Steps

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

```
storage aggregate relocation start -node source_node -destination destination-  
node -aggregate-list *
```

2. 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
```

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.
```


After you finish

You must restart the controller replacement process:

```
system controller replace start -nodes node_names
```

License

Some features require 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.

Refer to [References](#) to link to the *NetApp Support Site* where you can obtain new 28-character license keys for ONTAP. 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, refer to [References](#) to link to the *System Administration Reference*.

Manage storage encryption using the Onboard Key Manager

You can use the Onboard Key Manager (OKM) to manage encryption keys. If you have the OKM set up, you must record the passphrase and backup material before beginning the upgrade.

Steps

1. Record the cluster-wide passphrase.

This is the passphrase that was entered when the OKM was configured or updated using the CLI or REST API.

2. Back up the key-manager information by running the `security key-manager onboard show-backup` command.

Quiesce the SnapMirror relationships (optional)

Before continuing with the procedure, you must confirm that all the SnapMirror relationships are quiesced. When a SnapMirror relationship is quiesced, it remains quiesced across reboots and failovers.

Steps

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

```
snapmirror show
```



If the status is "Transferring", you must abort those transfers:

```
snapmirror abort -destination-vserver vserver_name
```

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

2. Quiesce all relationships between the cluster:

```
snapmirror quiesce -destination-vserver *
```

Stage 2. Relocate and retire node1

Relocate non-root aggregates and NAS data LIFs owned by node1 to node2

Before you can replace node1 with node3, you must move the non-root aggregates and NAS data LIFs from node1 to node2 before eventually moving node1's resources to node3.

Before you begin

The operation should already be paused when you begin the task; you must manually resume the operation.

About this task

After the aggregates and LIFs are migrated, the operation is paused for verification purposes. At this stage, you must verify whether or not all the non-root aggregates and non-SAN data LIFs are migrated to node3.



The home owner for the aggregates and LIFs is not modified; only the current owner is modified.

Steps

1. Resume the aggregate relocation and NAS data LIF move operations:

```
system controller replace resume
```

All the non-root aggregates and NAS data LIFs are migrated from node1 to node2.

The operation pauses to enable you to verify whether all node1 non-root aggregates and non-SAN data LIFs have been migrated to node2.

2. Check the status of the aggregate relocation and NAS data LIF move operations:

```
system controller replace show-details
```

3. With the operation still paused, verify that all the non-root aggregates are online for their state on node2:

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

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,normal							
aggr_2	825.0GB	825.0GB	0%	online	1	node2	
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 <aggregate_name>
```

4. Verify that all the volumes are online on node2 by using 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 <vserver_name> -volume <volume_name>
```

The `vserver_name` to use with this command is found in the output of the previous `volume show` command.

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

```
network interface modify -vserver vserver_name -lif LIF_name -home-node  
nodename -status-admin up
```

Relocate failed or vetoed aggregates to node2

If any aggregates fail to relocate or are vetoed, you need to manually relocate the aggregates to node2, or if necessary, override either the vetoes or destination checks.

About this task

The system pauses the relocation operation due to the error.

Steps

1. Check the event management system (EMS) logs to determine why the aggregate failed to relocate or was vetoed.
2. Relocate any failed or vetoed aggregates:

```
storage aggregate relocation start -node <node1> -destination <node2>
-aggregate-list <aggregate_name> -ndo-controller-upgrade true
```

3. When prompted, enter `y`.
4. You can force relocation by using one of the following methods:

Option	Description
Overriding veto checks	Use the following command: storage aggregate relocation start -node node1 -destination node2 -aggregate-list <aggregate_list> -ndo-controller-upgrade true -override-vetoes true
Overriding destination checks	Use the following command: storage aggregate relocation start -node node1 -destination node2 -aggregate-list <aggregate_list> -ndo-controller-upgrade true -override-vetoes true -override-destination-checks true

Retire node1

To retire node1, you resume the automated operation to disable the HA pair with node2 and shut node1 down correctly. Later in the procedure, you remove node1 from the rack or chassis.

Steps

1. Resume the operation:

```
system controller replace resume
```

2. Verify that node1 has been halted:

```
system controller replace show-details
```

After you finish

You can decommission node1 after the upgrade is completed. See [Decommission the old system](#).

Prepare for netboot

After you physically rack node3 and node4 later in the procedure, you might need to netboot them. The term "netboot" means you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you put a copy of the ONTAP 9 boot image onto a web server that the system can access.

Before you begin



- Verify that you can access a HTTP server with the system.
- Refer to [References](#) to link to the *NetApp Support Site* and download the necessary system files for your platform and the correct version of ONTAP.

About this task

You must netboot the new controllers if they do not have the same 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.

Steps

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

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

You will use the information in the directories in [Stage 3](#).

Stage 3. Install and boot node3

Install and boot node3

You must install node3 in the rack, transfer node1’s connections to node3, boot node3, and install ONTAP. You must then reassign any of node1’s spare disks, any disks belonging to the root volume, and any non-root aggregates that were not relocated to node2 earlier in the process, as outlined in this section.

About this task

The relocation operation is paused at the beginning of this stage. This process is largely automated; the operation pauses to enable you to check its status. You must manually resume the operation. In addition, you must verify the SAN LIFs have successfully moved to node3.

You need to netboot node3 if it does not have the same version of ONTAP 9 that is installed on node1. After you install node3, boot it 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, by following the instructions in [Prepare for netboot](#).



- For an AFF A800 or AFF C800 controller upgrade, you must ensure that all drives in the chassis are firmly seated against the midplane before removing node1. For more information, see [Replace the AFF A800 or AFF C800 controller modules](#).
- If you are upgrading a system with storage disks, you need to complete this entire section and then go to the [Configure FC ports on node3](#) and [Check and configure UTA/UTA2 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 the *Installation and Setup Instructions* for your node model.



If you're upgrading to a system with both nodes in the same chassis, install node4 and node3 in the chassis. If you don't install both nodes in the same chassis, when you boot node3, it behaves as if it were in a dual-chassis configuration, and when you boot node4, the interconnect between the nodes doesn't come up.

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

Cable the following connections using the *Installation and Setup Instructions* for the node3 platform, the appropriate disk shelf document, and the *HA pair management* documentation.

Refer to [References](#) to link to *HA pair management*.

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



You might 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 FC-VI cable connections from node1 to node3. If the new host does not have an FC-VI card, you might need to move the FC-VI card.

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 later.



When you boot node3, you might see the following warning 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.



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




Refer to [Prepare for netboot](#).

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



You must use the management port and IP as the netboot connection. Don't 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 using the following command at the boot environment prompt: <code>ifconfig e0M -auto</code>
Not running	<p>Manually configure the connection by using the following command at the boot environment prompt:</p> <pre>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></pre> <p><i>filer_addr</i> is the IP address of the storage system (mandatory). <i>netmask</i> is the network mask of the storage system (mandatory). <i>gateway</i> is the gateway for the storage system (mandatory). <i>dns_addr</i> is the IP address of a name server on your network (optional). <i>dns_domain</i> is the Domain Name Service (DNS) domain name (optional).</p> <div>  <p>Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p> </div>

7. Perform netboot on node3:

For...	Then...
FAS/AFF8000 series systems	<code>netboot http://<web_server_ip/path_to_web-accessible_directory>/netboot/kernel</code>
All other systems	<code>netboot http://<web_server_ip/path_to_web-accessible_directory>/<ontap_version>_image.tgz</code>

The `<path_to_the_web-accessible_directory>` should lead to where you downloaded the `<ontap_version>_image.tgz` in the section [Prepare for netboot](#).



Do not interrupt the boot.

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

This menu option downloads and installs the new 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 ONTAP, and not upgrades of controllers.



Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the incorrect image might install. This issue applies to all ONTAP releases. The netboot procedure combined with option (7) Install new software wipes the boot media and places the same ONTAP version on both image partitions.

9. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL:

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

10. Complete the following substeps to reboot the controller module:

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}
```

b. Enter `y` to reboot 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 must be restored.

11. Select maintenance mode 5 from the boot menu and enter `y` when you are prompted to continue with the boot.

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

```
ha-config show
```

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

```
Chassis HA configuration: ha
Controller HA configuration: ha
```



System records 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.

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

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

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

If you have a MetroCluster configuration, use the following commands to modify the controller and chassis:

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

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

14. Exit maintenance mode:

```
halt
```

Interrupt AUTOBOOT by pressing Ctrl-C at the boot environment prompt.

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

```
date
```

16. On node3, check the date by using the following command at the boot environment prompt:

```
show date
```

17. If necessary, set the date on node3:

```
set date mm/dd/yyyy
```

18. On node3, check the time by using the following command at the boot environment prompt:

```
show time
```

19. If necessary, set the time on node3:

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

20. In boot loader, set the partner system ID on node3:

```
setenv partner-sysid node2_sysid
```

For node3, partner-sysid must be that of node2.


a. Save the settings:

```
saveenv
```

21. Verify the partner-sysid for node3:

```
printenv partner-sysid
```


22. If you have NetApp Storage Encryption (NSE) drives installed, perform the following steps:



If you have not already done so earlier in the procedure, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use.

a. Set bootarg.storageencryption.support to true or false:

If the following drives are in use...	Then...
NSE drives that conform to FIPS 140-2 Level 2 self-encryption requirements	setenv bootarg.storageencryption.support true
NetApp non-FIPS SEDs	setenv bootarg.storageencryption.support false



You cannot mix FIPS drives with other types of drives on the same node or HA pair. You can mix SEDs with non-encrypting drives on the same node or HA pair.

b. Go to the special boot menu and select option (10) Set Onboard Key Manager recovery secrets.

Enter the passphrase and the backup information that you recorded earlier procedure. See [Manage storage encryption using the Onboard Key Manager](#).

23. Boot the node into boot menu:

```
boot_ontap menu
```

What's next?

- If you have a system with an FC or UTA/UTA2 configuration, [set and configure the FC or UTA/UTA2 ports on node3](#).
- If you don't have an FC or UTA/UTA2 configuration, [reassign node1 disks to node3, Step 1](#) so that node3 can recognize node1's disks.
- If you have a MetroCluster configuration, [set and configure the FC or UTA/UTA2 ports on node3](#) to detect the disks attached to the node.

Set the FC or UTA/UTA2 configuration on node3

If node3 has onboard FC ports, onboard unified target adapter (UTA/UTA2) ports, or a UTA/UTA2 card, you must configure the settings before completing the rest of the procedure.

About this task

You might need to complete the section [Configure FC ports on node3](#), the section [Check and configure UTA/UTA2 ports on node3](#), or both sections.



NetApp marketing materials might use the term UTA2 to refer to converged network adapter (CNA) adapters and ports. However, the CLI uses the term CNA.

If node3 does not have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card (for example, AFF and FAS systems introduced beginning with ONTAP 9.15.1), and you are upgrading a system with storage disks, you can skip to [Reassign node1 disks to node3](#).

Configure FC ports on node3

If node3 has FC ports, either onboard or on an add-on FC adapter, you must set port configurations on the node before you bring it into service because the ports are not preconfigured when the systems are shipped. If you don't configure the ports, 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 [Prepare the nodes for upgrade](#).

About this task

You can skip this section if your system does not have FC configurations. If your system has onboard UTA/UTA2 ports or a UTA/UTA2 card, you configure them in [Check and configure UTA/UTA2 ports on node3](#).



Enter the commands in this section at the Maintenance mode shell prompt.

Steps

1. Compare the FC settings on node3 with the settings that you captured earlier from node1.
2. Take one of the following actions to modify the FC ports on node3, as needed:

In Maintenance mode (option 5 at boot menu):

- To program as target ports:

```
ucadmin modify -m fc -t target <adapter>
```

For example: `ucadmin modify -m fc -t target 2a`

- To program initiator ports:

```
ucadmin modify -m fc -t initiator <adapter>
```

For example: `ucadmin modify -m fc -t initiator 2b`

3. Verify the new settings by using the following command and examining the output:

```
ucadmin show
```

4. Halt the node:

```
halt
```

5. Boot the system from LOADER prompt:

```
boot_ontap menu
```

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

7. Select option 5 from the boot menu for maintenance mode.

8. Take one of the following actions:

If node 3...	Then...
Has a UTA/UTA2 card or UTA/UTA2 onboard ports	Go to Check and configure UTA/UTA2 ports on node3
Does not have a UTA/UTA2 card or UTA/UTA2 onboard ports	Skip <i>Check and configure UTA/UTA2 ports on node3</i> and go to Reassign node1 disks to node3 .

Check and configure UTA/UTA2 ports on node3

If node3 has onboard UTA/UTA2 ports or a UTA/UTA2 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 UTA/UTA2 ports.

About this task

If you want to use a Unified Target Adapter (UTA/UTA2) port for FC, you must first verify how the port is configured.



NetApp marketing materials might use the term UTA2 to refer to CNA adapters and ports. However, the CLI uses the term CNA.

You can use the `ucadmin show` command to view or verify the current port configuration, as shown in the following example output:

```
*> ucadmin show
```

Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
0e	fc	target	-	initiator	offline
0f	fc	target	-	initiator	offline
0g	fc	target	-	initiator	offline
0h	fc	target	-	initiator	offline
1a	fc	target	-	-	online
1b	fc	target	-	-	online

6 entries were displayed.

UTA/UTA2 ports can be configured into native FC mode or UTA/UTA2 mode. FC mode supports FC initiator and FC target; UTA/UTA2 mode allows concurrent NIC and FCoE traffic sharing the same 10GbE SFP+ interface and supports FC targets.

You might find UTA/UTA2 ports on an add-on adapter or on the controller motherboard, and have the following configurations, but you should check the configuration of the UTA/UTA2 ports on the node3 and change it, if necessary:

- UTA/UTA2 cards ordered when the controller is ordered are configured before shipment to have the personality you request.
- UTA/UTA2 cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard UTA/UTA2 ports on new controllers are configured before shipment to have the personality you request.



You must be in Maintenance mode to configure UTA/UTA2 ports. Enter the commands in this section at the Maintenance mode shell prompt.

Steps

1. 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.

2. Verify the UTA/UTA2 port settings:

```
ucadmin show
```

Examine the output and determine whether the UTA/UTA2 ports have the personality you want.

The output in the following example shows that the type of adapter "1b" is changing to initiator and that the mode of adapters "2a" and "2b" is changing to "cna". The CNA mode allows you to use the card as a network adapter.

```
*> 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

```
*>
```

3. Take one of the following actions:

If the UTA/UTA2 ports...	Then...
Do not have the personality that you want	Go to Step 4 .
Have the personality that you want	Skip Step 4 through Step 8 and go to Step 9 .

4. Take one of the following actions:

If you are configuring...	Then...
Ports on a UTA/UTA2 card	Go to Step 5
Onboard UTA/UTA2 ports	Skip Step 5 and go to Step 6 .

5. If the adapter is in initiator mode, and if the UTA/UTA2 port is online, take the UTA/UTA2 port offline:

```
storage disable adapter <adapter_name>
```

Adapters in target mode are automatically offline in Maintenance mode.

6. 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 is the FC4 type, target or initiator.



You must use FC initiator for tape drives and MetroCluster configurations. You must use the FC target for SAN clients.

7. Place any target ports online by entering the following command once for each port:

```
storage enable adapter <adapter_name>
```

8. Cable the port.

9. Exit maintenance mode:

```
halt
```

10. Boot the node into boot menu:

```
boot_ontap menu
```

What's next?

- If you are upgrading to an AFF A800 system, go to [Reassign node1 disks to node3, Step 9](#).
- For all other system upgrades, go to [Reassign node1 disks to node3, Step 1](#).

Reassign node1 disks to node3

You need to reassign the disks that belonged to node1 to node3 before verifying the node3 installation.

Steps

1. Verify that node1 has stopped at the boot menu. Reassign the disks of node1 to node3:

```
boot_after_controller_replacement
```

After a short delay, you are prompted to enter the name of the node that is being replaced. If there are shared disks (also called Advanced Disk Partitioning (ADP) or partitioned disks), you are prompted to enter the node name of the HA partner.

These prompts might get buried in the console messages. If you do not enter a node name or enter an incorrect name, you are prompted to enter the name again.

Expand the console output example

```
LOADER-A> boot_ontap menu
.
<output truncated>
.
All rights reserved.
*****
*                                     *
* Press Ctrl-C for Boot Menu. *
*                                     *
*****
.
<output truncated>
.
Please choose one of the following:
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 22/7
(22/7) Print this secret List
(25/6) Force boot with multiple filesystem disks missing.
(25/7) Boot w/ disk labels forced to clean.
(29/7) Bypass media errors.
(44/4a) Zero disks if needed and create new flexible root volume.
(44/7) Assign all disks, Initialize all disks as SPARE, write DDR
labels
.
<output truncated>
.
(wipeconfig) Clean all configuration on boot
device
(boot_after_controller_replacement) Boot after controller upgrade
(boot_after_mcc_transition) Boot after MCC transition
(9a) Unpartition all disks and remove
their ownership information.
(9b) Clean configuration and
initialize node with partitioned disks.
```



```

(9c)                                Clean configuration and
initialize node with whole disks.
(9d)                                Reboot the node.
(9e)                                Return to main boot menu.
The boot device has changed. System configuration information could
be lost. Use option (6) to restore the system configuration, or
option (4) to initialize all disks and setup a new system.
Normal Boot is prohibited.
Please choose one of the following:
(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? boot_after_controller_replacement
This will replace all flash-based configuration with the last backup
to disks. Are you sure you want to continue?: yes
.
<output truncated>
.
Controller Replacement: Provide name of the node you would like to
replace:<nodename of the node being replaced>
Controller Replacement: Provide High Availability partner of node1:
<nodename of the partner of the node being replaced>
Changing sysid of node node1 disks.
Fetched sanown old_owner_sysid = 536940063 and calculated old sys id
= 536940063
Partner sysid = 4294967295, owner sysid = 536940063
.
<output truncated>
.
varfs_backup_restore: restore using /mroot/etc/varfs.tgz
varfs_backup_restore: attempting to restore /var/kmip to the boot
device
varfs_backup_restore: failed to restore /var/kmip to the boot device
varfs_backup_restore: attempting to restore env file to the boot
device
varfs_backup_restore: successfully restored env file to the boot
device wrote key file "/tmp/rndc.key"
varfs_backup_restore: timeout waiting for login

```

```

varfs_backup_restore: Rebooting to load the new varfs
Terminated
<node reboots>
System rebooting...
.
Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...
Rebooting to load the restored env file...
.
System rebooting...
.
<output truncated>
.
WARNING: System ID mismatch. This usually occurs when replacing a
boot device or NVRAM cards!
Override system ID? {y|n} y
.
Login:

```



In the above console output example, ONTAP will prompt you for the partner node name if the system uses Advanced Disk Partitioning (ADP) disks.

2. If the system goes into a reboot loop with the message `no disks found`, it indicates that the system has reset the FC or UTA/UTA2 ports back to the target mode and therefore is unable to see any disks. Select one of the following tasks to resolve this issue:

- Perform [Step 3](#) to [Step 8](#) on node3
- Go to section [Verify the node3 installation](#)

3. Press Ctrl-C during AUTOBOOT to stop the node at the `LOADER>` prompt.

4. At the `LOADER` prompt, enter maintenance mode:

```
boot_ontap maint
```

5. In maintenance mode, display all the previously set initiator ports that are now in target mode:

```
ucadmin show
```

Change the ports back to initiator mode:

```
ucadmin modify -m fc -t initiator -f adapter name
```

6. Verify that the ports have been changed to initiator mode:

```
ucadmin show
```

7. Exit maintenance mode:

halt



If you are upgrading from a system that supports external disks to a system that also supports external disks, go to [Step 8](#).

If you are upgrading from a system that supports external disks to a system that supports both internal and external disks, for example, an AFF A800 system, go to [Step 9](#).

8. At the LOADER prompt, boot up:

```
boot_ontap menu
```

Now, on booting, the node can detect all the disks that were previously assigned to it and can boot up as expected.

When the cluster nodes you are replacing use root volume encryption, ONTAP is unable to read the volume information from the disks. Restore the keys for the root volume.



This only applies when the root volume is using NetApp Volume Encryption.

a. Return to the special boot menu:

```
LOADER> boot_ontap menu
```

```
Please choose one of the following:
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.

Selection (1-11)? 10
```

b. Select **(10) Set Onboard Key Manager recovery secrets**

c. Enter **y** at the following prompt:

```
This option must be used only in disaster recovery procedures. Are you sure?
(y or n): y
```

d. At the prompt, enter the key-manager passphrase.

e. Enter the backup data when prompted.



You must have obtained the passphrase and backup data in the [Prepare the nodes for upgrade](#) section of this procedure.

- f. After the system boots to the special boot menu again, run option **(1) Normal Boot**



You might encounter an error at this stage. If an error occurs, repeat the substeps in [Step 8](#) until the system boots normally.

9. If you are upgrading from a system with external disks to a system that supports internal and external disks (AFF A800 systems, for example), set the node1 aggregate as the root aggregate to confirm that node3 boots from the root aggregate of node1. To set the root aggregate, go to the boot menu and select option 5 to enter maintenance mode.



You must perform the following substeps in the exact order shown; failure to do so might cause an outage or even data loss.

The following procedure sets node3 to boot from the root aggregate of node1:

- a. Enter maintenance mode:

```
boot_ontap maint
```

- b. Check the RAID, plex, and checksum information for the node1 aggregate:

```
aggr status -r
```

- c. Check the status of the node1 aggregate:

```
aggr status
```

- d. If necessary, bring the node1 aggregate online:

```
aggr_online root_aggr_from_node1
```

- e. Prevent the node3 from booting from its original root aggregate:

```
aggr offline root_aggr_on_node3
```

- f. Set the node1 root aggregate as the new root aggregate for node3:

```
aggr options aggr_from_node1 root
```

- g. Verify that the root aggregate of node3 is offline and the root aggregate for the disks brought over from node1 is online and set to root:

```
aggr status
```



Failing to perform the previous substep might cause node3 to boot from the internal root aggregate, or it might cause the system to assume a new cluster configuration exists or prompt you to identify one.

The following shows an example of the command output:

Aggr	State	Status	Options
aggr0_nst_fas8080_15	online	raid_dp, aggr fast zeroed 64-bit	root, nosnap=on
aggr0	offline	raid_dp, aggr fast zeroed 64-bit	diskroot

Verify the node3 installation

You must verify that the physical ports from node1 map correctly to the physical ports on node3. This will enable node3 to communicate with other nodes in the cluster and with the network after the upgrade.

About this task

Refer to [References](#) to link to the *Hardware Universe* to capture information about the ports on the new nodes. You will use the information later in this section.

Physical port layout might vary, depending on the model of the nodes. When the new node boots up, ONTAP will try to determine which ports should host cluster LIFs in order to automatically come into quorum.

If the physical ports on node1 do not map directly to the physical ports on node3, the subsequent section [Restore network configuration on node3](#) must be used to repair the network connectivity.

After you install and boot node3, you must verify that it is installed correctly. You must wait for node3 to join quorum and then resume the relocation operation.

At this point in the procedure, the operation will have paused as node3 joins quorum.

Steps

1. Verify that node3 has joined quorum:

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

The output of the `health` field should be `true`.

2. Verify that node3 is part of the same cluster as node2 and that it is healthy:

```
cluster show
```

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

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

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

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

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

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

Repeat Step 3 to verify that the cluster LIF is now listening on port 7700.

- Switch to advanced privilege mode:

```
set advanced
```

- Check the status of the controller replacement operation and verify that it is in a paused state and in the same state that it was in before node1 was halted to perform the physical tasks of installing new controllers and moving cables:

```
system controller replace show
```

```
system controller replace show-details
```

- If you are working on a MetroCluster system, verify that the replaced controller is configured correctly for the MetroCluster configuration; the MetroCluster configuration should be in a healthy state. See [Verify the health of the MetroCluster configuration](#).

Reconfigure the intercluster LIFs on MetroCluster node node3, and check cluster peering to restore communication between the MetroCluster nodes before proceeding to Step 6.

Check the MetroCluster node status:

```
metrocluster node show
```

8. Resume the controller replacement operation:

```
system controller replace resume
```

9. Controller replacement will pause for intervention with the following message:

```
Cluster::*> system controller replace show
Node           Status           Error-Action
-----
Node1(now node3) Paused-for-intervention Follow the instructions
given in
Step Details
Node2           None
Step Details:
-----
To complete the Network Reachability task, the ONTAP network
configuration must be manually adjusted to match the new physical
network configuration of the hardware. This includes:

1. Re-create the interface group, if needed, before restoring VLANs. For
detailed commands and instructions, refer to the "Re-creating VLANs,
ifgrps, and broadcast domains" section of the upgrade controller
hardware guide for the ONTAP version running on the new controllers.
2. Run the command "cluster controller-replacement network displaced-
vans show" to check if any VLAN is displaced.
3. If any VLAN is displaced, run the command "cluster controller-
replacement network displaced-vans restore" to restore the VLAN on the
desired port.

2 entries were displayed.
```



In this procedure, the section *Re-creating VLANs, ifgrps, and broadcast domains* has been renamed *Restore network configuration on node3*.

10. With the controller replacement in a paused state, proceed to the next section of this document to restore network configuration on the node.

Restore network configuration on node3

After you confirm that node3 is in quorum and can communicate with node2, verify that node1's VLANs, interface groups, and broadcast domains are seen on node3. Also, verify that all node3 network ports are configured in their correct broadcast domains.

About this task

For more information on creating and re-creating VLANs, interface groups, and broadcast domains, refer to [References](#) to link to *Network Management*.



If you are changing the port speed of the e0a and e1a cluster ports on AFF A800 or AFF C800 systems, you might observe malformed packets being received after the speed conversion. See [NetApp Bugs Online Bug ID 1570339](#) and the knowledge base article [CRC errors on T6 ports after converting from 40GbE to 100GbE](#) for guidance.

Steps

1. List all the physical ports that are on upgraded node1 (referred to as node3):

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

All physical network ports, VLAN ports and interface group ports on the node are displayed. From this output, you can see any physical ports that have been moved into the `Cluster` broadcast domain by ONTAP. You can use this output to aid in deciding which ports must be used as interface group member ports, VLAN base ports or standalone physical ports for hosting LIFs.

2. List the broadcast domains on the cluster:

```
network port broadcast-domain show
```

3. List the network port reachability of all ports on node3:

```
network port reachability show
```

You should see output like the following example:


```
clusterA::*> reachability show -node node1_node3
(network port reachability show)
Node          Port          Expected Reachability  Reachability Status
-----
node1_node3
a0a           Default:Default        no-reachability
a0a-822       Default:822            no-reachability
a0a-823       Default:823            no-reachability
e0M           Default:Mgmt           ok
e0a           Cluster:Cluster        misconfigured-
reachability
e0b           Cluster:Cluster        no-reachability
e0c           Cluster:Cluster        no-reachability
e0d           Cluster:Cluster        no-reachability
e0e           Cluster:Cluster        ok
e0e-822       -                      no-reachability
e0e-823       -                      no-reachability
e0f           Default:Default        no-reachability
e0f-822       Default:822            no-reachability
e0f-823       Default:823            no-reachability
e0g           Default:Default        misconfigured-
reachability
e0h           Default:Default        ok
e0h-822       Default:822            ok
e0h-823       Default:823            ok
18 entries were displayed.
```

In the preceding example, node1_node3 is just booted after controller replacement. Some ports do not have reachability to their expected broadcast domains and must be repaired.

4. Repair the reachability for each of the ports on node3 with a reachability status other than `ok`. Run the following command, first on any physical ports, then on any VLAN ports, one at a time:

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

You should see output like the following example:

```
Cluster ::> reachability repair -node node1_node3 -port e0h
```

```
Warning: Repairing port "node1_node3: e0h" may cause it to move into a
different broadcast domain, which can cause LIFs to be re-homed away
from the port. Are you sure you want to continue? {y|n}:
```

A warning message, as shown above, is expected for ports with a reachability status that might be different

from the reachability status of the broadcast domain where it is currently located. Review the connectivity of the port and answer *y* or *n* as appropriate.

Verify that all physical ports have their expected reachability:

```
network port reachability show
```

As the reachability repair is performed, ONTAP attempts to place the ports in the correct broadcast domains. However, if a port's reachability cannot be determined and does not belong to any of the existing broadcast domains, ONTAP will create new broadcast domains for these ports.

5. If interface group configuration does not match the new controller physical port layout, modify it by using the following steps.
 - a. You must first remove physical ports that should be interface group member ports from their broadcast domain membership. You can do this by using the following command:

```
network port broadcast-domain remove-ports -broadcast-domain broadcast-domain_name -ports node_name:port_name
```

- b. Add a member port to an interface group:

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

- c. The interface group is automatically added to the broadcast domain about a minute after the first member port is added.
 - d. Verify that the interface group was added to the appropriate broadcast domain:

```
network port reachability show -node node_name -port ifgrp
```

If the interface group's reachability status is not *ok*, assign it to the appropriate broadcast domain:

```
network port broadcast-domain add-ports -broadcast-domain broadcast_domain_name -ports node:port
```

6. Assign appropriate physical ports to the *Cluster* broadcast domain by using the following steps:

- a. Determine which ports have reachability to the *Cluster* broadcast domain :

```
network port reachability show -reachable-broadcast-domains Cluster:Cluster
```

- b. Repair any port with reachability to the *Cluster* broadcast domain, if its reachability status is not *ok*:

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

7. Move the remaining physical ports into their correct broadcast domains by using one of the following commands:

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

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

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

Verify that there are no unreachable or unexpected ports present. Check the reachability status for all physical ports by using the following command and examining the output to confirm the status is ok:

```
network port reachability show -detail
```

8. Restore any VLANs that might have become displaced by using the following steps:

a. List displaced VLANs:

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

Output like the following should display:

```
Cluster::*> displaced-vlans show
(cluster controller-replacement network displaced-vlans show)
      Original
Node   Base Port   VLANs
-----
Node1   a0a           822, 823
        e0e           822, 823
2 entries were displayed.
```

b. Restore VLANs that were displaced from their previous base ports:

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

The following is an example of restoring VLANs that have been displaced from interface group "a0a" back onto the same interface group:

```
Cluster::*> displaced-vlans restore -node node1_node3 -port a0a
-destination-port a0a
```

The following is an example of restoring displaced VLANs on port "e0e" to 'e0h':

```
Cluster::*> displaced-vlans restore -node node1_node3 -port e0e
-destination-port e0h
```

When a VLAN restore is successful, the displaced VLANs are created on the specified destination port. The VLAN restore fails if the destination port is a member of an interface group, or if the destination port is down.

Wait about one minute for newly restored VLANs to be placed into their appropriate broadcast domains.

c. Create new VLAN ports as needed for VLAN ports that are not in the `cluster controller-replacement network displaced-vlans show` output but should be configured on other physical ports.

9. Delete any empty broadcast domains after all port repairs have been completed:

```
network port broadcast-domain delete -broadcast-domain broadcast_domain_name
```

10. Verify port reachability:

```
network port reachability show
```

When all ports are correctly configured and added to the correct broadcast domains, the `network port reachability show` command should report the reachability status as `ok` for all connected ports, and the status as `no-reachability` for ports with no physical connectivity. If any port reports a status other than these two, perform the reachability repair and add or remove ports from their broadcast domains as instructed in [Step 4](#).

11. Verify that all ports have been placed into broadcast domains:

```
network port show
```

12. Verify that all ports in the broadcast domains have the correct maximum transmission unit (MTU) configured:

```
network port broadcast-domain show
```

13. Restore LIF home ports, specifying the Vserver(s) and LIF(s) home ports, if any, that need to be restored by using the following steps:

- a. List any LIFs that are displaced:

```
displaced-interface show
```

- b. Restore LIF home nodes and home ports:

```
cluster controller-replacement network displaced-interface restore-home-node  
-node node_name -vserver vserver_name -lif-name LIF_name
```

14. Verify that all LIFs have a home port and are administratively up:

```
network interface show -fields home-port, status-admin
```

Restore key-manager configuration on node3

If you are using NetApp Volume Encryption (NVE) and NetApp Aggregate Encryption (NAE) to encrypt volumes on the system you are upgrading, the encryption configuration must be synchronized to the new nodes. If you do not synchronize the key-manager, when you relocate the node1 aggregates from node2 to node3 by using ARL, failures might occur because node3 does not have the required encryption keys to bring encrypted volumes and aggregates online.

About this task

Synchronize the encryption configuration to the new nodes by performing the following steps:

Steps

1. Run the following command from node3:

```
security key-manager onboard sync
```

2. Verify that the SVM-KEK key is restored to "true" on node3 before you relocate the data aggregates:

```
::> security key-manager key query -node node3 -fields restored -key  
-type SVM-KEK
```

Example

```
::> security key-manager key query -node node3 -fields restored -key  
-type SVM-KEK
```

node	vserver	key-server	key-id
restored			
-----	-----	-----	-----

node3	svm1	" "	0000000000000000020000000000a008a81976
true			2190178f9350e071fbb90f00000000000000000

Move non-root aggregates and NAS data LIFs owned by node1 from node2 to node3

After you verify network configuration on node3 and before you relocate aggregates from node2 to node3, you must verify that the NAS data LIFs belonging to node1 that are currently on node2 are relocated from node2 to node3. You must also verify that the SAN LIFs exist 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.



If you are changing the port speed of the T6-based Ethernet network interface cards or motherboard ports, you might observe malformed packets being received after the speed conversion. See [NetApp Bugs Online Bug ID 1570339](#) and the knowledge base article [CRC errors on T6 ports after converting from 40GbE to 100GbE](#) for guidance.

Steps

1. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

- Cluster quorum check
- System ID check
- Image version check
- Target platform check
- Network reachability check

The operation pauses at this stage in the network reachability check.

2. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following checks:

- Cluster health check
- Cluster LIF status check

After performing these checks, the system relocates the non-root aggregates and NAS data LIFs owned by node1 to the new controller, node3.

The controller replacement operation pauses after the resource relocation is complete.

3. Check the status of the aggregate relocation and NAS data LIF move operations:

```
system controller replace show-details
```

If the controller replacement procedure is paused, check and correct the error, if any, and then issue `resume` to continue the operation.

4. If necessary, restore and revert displaced LIFs or manually migrate and modify the node1 LIFs that failed to relocate automatically to node3.

Restore and revert displaced LIFs

- a. List any displaced LIFs:

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

- b. If any LIFs are displaced, restore the home node back to node3:

```
cluster controller-replacement network displaced-interface  
restore-home-node -node <node3_nodename> -vserver <vserver name>  
-lif-name <lif_name>
```

Manually migrate and modify LIFs

- a. Migrate the LIFs that failed to relocate automatically to node3:

```
network interface migrate -vserver <vserver name> -lif <lif_name>  
-destination-node <node3_nodename> -destination-port  
<port_on_node3>
```

- b. Modify the home node and home port for the migrated LIFs:

```
network interface modify -vserver <vserver_name> -lif  
<data_lif_name> -home-node <node3_nodename> -home-port  
<home_port>
```

5. Resume the operation to prompt the system to perform the required post-checks:

```
system controller replace resume
```

The system performs the following post-checks:

- Cluster quorum check
- Cluster health check
- Aggregates reconstruction check
- Aggregate status check
- Disk status check
- Cluster LIF status check
- Volume check

Stage 4. Relocate and retire node2

Relocate non-root aggregates and NAS data LIFs from node2 to node3

Before replacing node2 with node4, you relocate the non-root aggregates and NAS data LIFs that are owned by node2 to node3.

Before you begin

After the post-checks from the previous stage complete, the resource release for node2 starts automatically. The non-root aggregates and non-SAN data LIFs are migrated from node2 to 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.

After the aggregates and LIFs are migrated, the operation is paused for verification purposes. At this stage, you must verify whether or not all the non-root aggregates and non-SAN data LIFs are migrated to node3.



The home owner for the aggregates and LIFs are not modified; only the current owner is modified.

Steps

1. Verify that all the non-root aggregates are online and their state on node3:

```
storage aggregate show -node <node3> -state online -root false
```

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

```
cluster::> storage aggregate show -node node3 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	normal					
aggr_2	825.0GB	825.0GB	0%	online	1	node2
raid_dp	normal					

2 entries were displayed.

If the aggregates have gone offline or become foreign on node3, bring them online by using the following command on node3, once for each aggregate:

```
storage aggregate online -aggregate <aggregate_name>
```

2. Verify that all the volumes are online on node3 by using the following command on node3 and examining the output:

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


If any volumes are offline on node3, bring them online by using the following command on node3, once for each volume:

```
volume online -vserver <vserver_name> -volume <volume_name>
```

The `vserver_name` to use with this command is found in the output of the previous `volume show` command.

3. Verify that the LIFs have been moved to the correct ports and have a status of `up`. 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> -home-node <node_name> -status-admin up
```

4. If the ports currently hosting data LIFs will not exist on the new hardware, remove them from the broadcast domain:

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

5. Verify that there are no data LIFs remaining on node2 by entering the following command and examining the output:

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

Relocate failed or vetoed aggregates to node3

If any aggregates fail to relocate or are vetoed, you need to manually relocate the aggregates to node3, or if necessary, override either the vetoes or destination checks.

About this task

The system pauses the relocation operation due to the error.

Steps

1. Check the event management system (EMS) logs to determine why the aggregate failed to relocate or was vetoed.
2. Relocate any failed or vetoed aggregates:

```
storage aggregate relocation start -node <node2> -destination <node3>  
-aggregate-list <aggregate_name> -ndo-controller-upgrade true
```

3. When prompted, enter `y`.
4. You can force relocation by using one of the following methods:

Option	Description
Overriding veto checks	Use the following command: storage aggregate relocation start -node node2 -destination node3 -aggregate-list <aggregate_list> -ndo-controller-upgrade true -override-vetoes true

Option	Description
Overriding destination checks	Use the following command: <pre>storage aggregate relocation start -node node2 -destination node3 -aggregate-list <aggregate_list> -ndo-controller-upgrade true -override-vetoes true -override-destination-checks true</pre>

Retire node2

To retire node2, you shut down node2 correctly and then remove it from the rack or chassis.

Steps

1. Resume the operation:

```
system controller replace resume
```

The node halts automatically.

After you finish

You can decommission node2 after the upgrade is completed. See [Decommission the old system](#).

Stage 5. Install and boot node4

Install and boot node4

You must install node4 in the rack, transfer node2's connections to node4, boot node4, and install ONTAP. You must then reassign any of node2's spare disks, any disks belonging to the root volume, and any non-root aggregates that were not relocated to node3 earlier in the process, as outlined in this section.

About this task

The relocation operation is paused at the beginning of this stage. This process is mostly automated; the operation pauses to enable you to check its status. You must manually resume the operation.

You need to netboot node4 if the ONTAP version on node4 is different to the ONTAP version on node2. After you install node4, boot it 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 by following the instructions in [Prepare for netboot](#).



- For an AFF A800 or AFF C800 controller upgrade, you must ensure that all drives in the chassis are firmly seated against the midplane before removing node2. For more information, see [Replace the AFF A800 or AFF C800 controller modules](#).
- If you are upgrading a system with storage disks, you must complete this entire section and then proceed to [Set the FC or UTA/UTA2 configuration on node4](#), entering commands at the cluster prompt.

Steps

1. Make sure that node4 has sufficient rack space.

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

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

Cable the following connections, using the *Installation and Setup Instructions* for the node4 platform, the appropriate disk shelf document, and the *HA pair management* documentation.

Refer to [References](#) to link to *HA pair management*.

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



You might 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.

For the MetroCluster configuration, you must move the FC-VI cable connections from node2 to node4. If the new host does not have an FC-VI card, you might need to move the FC-VI card.

4. 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.



When you boot node4, you might see the following warning 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.



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




Refer to [Prepare for netboot](#).

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



You must use the management port and IP as the netboot connection. Don't 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	<p>Configure the connection automatically by using the following command at the boot environment prompt:</p> <pre>ifconfig e0M -auto</pre>
Not running	<p>Manually configure the connection by entering the following command at the boot environment prompt:</p> <pre>ifconfig e0M -addr=filer_addr -mask=netmask -gw=gateway -dns=dns_addr -domain=dns_domain</pre> <p><i>filer_addr</i> is the IP address of the storage system (mandatory). <i>netmask</i> is the network mask of the storage system (mandatory). <i>gateway</i> is the gateway for the storage system (mandatory). <i>dns_addr</i> is the IP address of a name server on your network (optional). <i>dns_domain</i> is the DNS domain name (optional).</p> <div>  <p>Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p> </div>

- Perform netboot on node4:

For...	Then...
FAS/AFF8000 series systems	<code>netboot http://<web_server_ip/path_to_web-accessible_directory>/netboot/kernel</code>
All other systems	<code>netboot http://<web_server_ip/path_to_web-accessible_directory>/<ontap_version>_image.tgz</code>

The `<path_to_the_web-accessible_directory>` should lead to where you downloaded the `<ontap_version>_image.tgz` in [Step 1](#) in the section [Prepare for netboot](#).



Do not interrupt the boot.

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

This menu option downloads and installs the new 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 ONTAP, and not upgrades of controllers.



Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the incorrect image might install. This issue applies to all ONTAP releases. The netboot procedure combined with option (7) `Install new software` wipes the boot media and places the same ONTAP version on both image partitions.

9. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL:

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

10. Complete the following substeps to reboot the controller module:

- 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}
```

- 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 must be restored.

11. Select maintenance mode 5 from the boot menu and enter `y` when you are prompted to continue with the boot.
12. Verify that the controller and chassis are configured as HA:

```
ha-config show
```

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

```
Chassis HA configuration: ha  
Controller HA configuration: ha
```



System records 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.

13. If the controller and chassis are not configured as HA, use the following commands to correct the configuration:

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

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

If you have a MetroCluster configuration, use the following commands to modify the controller and chassis:

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

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

14. Exit maintenance mode:

```
halt
```

Interrupt AUTOBOOT by pressing Ctrl-C at the boot environment prompt.

15. On node3, check the system date, time, and time zone:

```
date
```

16. On node4, check the date by using the following command at the boot environment prompt:

```
show date
```

17. If necessary, set the date on node4:

```
set date mm/dd/yyyy
```

18. On node4, check the time by using the following command at the boot environment prompt:

```
show time
```

19. If necessary, set the time on node4:

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

20. In boot loader, set the partner system ID on node4:

```
setenv partner-sysid node3_sysid
```

For node4, partner-sysid must be that of node3.

Save the settings:

```
saveenv
```

21. Verify the partner-sysid for node4:

```
printenv partner-sysid
```

22. If you have NetApp Storage Encryption (NSE) drives installed, perform the following steps:



If you have not already done so earlier in the procedure, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use.

a. Set `bootarg.storageencryption.support` to `true` or `false`:

If the following drives are in use...	Then...
NSE drives that conform to FIPS 140-2 Level 2 self-encryption requirements	<code>setenv bootarg.storageencryption.support true</code>
NetApp non-FIPS SEDs	<code>setenv bootarg.storageencryption.support false</code>



You cannot mix FIPS drives with other types of drives on the same node or HA pair. You can mix SEDs with non-encrypting drives on the same node or HA pair.

b. Go to the special boot menu and select option (10) `Set Onboard Key Manager recovery secrets`.

Enter the passphrase and the backup information that you recorded earlier procedure. See [Manage storage encryption using the Onboard Key Manager](#).

23. Boot the node into boot menu:

```
boot_ontap menu
```

What's next?

- If you have an FC or UTA/UTA2 configuration, [set and configure the FC or UTA/UTA2 ports on node4](#).
- If you don't have an FC or UTA/UTA2 configuration, [reassign node2 disks to node4, Step 1](#) so that node4 can recognize node2's disks.
- If you have a MetroCluster configuration, [set and configure the FC or UTA/UTA2 ports on node4](#) to detect the disks attached to the node.

Set the FC or UTA/UTA2 configuration on node4

If node4 has onboard FC ports, onboard unified target adapter (UTA/UTA2) ports, or a UTA/UTA2 card, you must configure the settings before completing the rest of the procedure.

About this task

You might need to complete [Configure FC ports on node4](#) or [Check and configure UTA/UTA2 ports on node4](#), or both sections.



If node4 does not have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card (for example, AFF and FAS systems introduced beginning with ONTAP 9.15.1), and you are upgrading a system with storage disks, you can skip to [Reassign node2 disks to node4](#).

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

Configure FC ports on node4

If node4 has FC ports, either onboard or on an add-on FC adapter, you must set port configurations on the node before you bring it into service because the ports are not preconfigured when the systems are shipped. If you don't configure the ports as required, 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 [Prepare the nodes for upgrade](#).

About this task

You can skip this section if your system does not have FC configurations. If your system has onboard UTA/UTA2 ports or a UTA/UTA2 adapter, you configure them in [Check and configure UTA/UTA2 ports on node4](#).



Enter the commands in this section at the Maintenance mode shell prompt.

Steps

1. Display information about all FC and converged network adapters on the system:

```
system node hardware unified-connect show
```

2. Compare the FC settings on node4 with the settings that you captured earlier from node1.
3. Modify the FC ports on node4 as needed:

- To program as target ports:

```
ucadmin modify -m fc -t target adapter
```

For example: `ucadmin modify -m fc -t target 2a`

- To program initiator ports:

```
ucadmin modify -m fc -t initiator adapter
```

`-t` is the FC4 type: target or initiator.

For example: `ucadmin modify -m fc -t initiator 2b`

4. Halt the node:

```
halt
```

5. Boot the system from LOADER prompt:


```
boot_ontap menu
```

6. After you enter the command, wait until the system stops at the boot environment prompt.
7. Select option 5 from the boot menu for maintenance mode.
8. Take one of the following actions:
 - Go to [Check and configure UTA/UTA2 ports on node4](#) if node4 has a UTA/UTA2 card or UTA/UTA2 onboard ports.
 - If node4 doesn't have a UTA/UTA2 card or UTA/UTA2 onboard ports, skip *Check and configure UTA/UTA2 ports on node4* and go to [Reassign node2 disks to node4](#).

Check and configure UTA/UTA2 ports on node4

If node4 has onboard UTA/UTA2 ports or a UTA/UTA2A 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 UTA/UTA2 ports.

About this task

UTA/UTA2 ports can be configured into native FC mode or UTA/UTA2A mode. FC mode supports FC initiator and FC target; UTA/UTA2 mode allows concurrent NIC and FCoE traffic to share the same 10GbE SFP+ interface and supports FC target.



NetApp marketing materials might use the term UTA2 to refer to CNA adapters and ports. However, the CLI uses the term CNA.

UTA/UTA2 ports might be on an adapter or on the controller with the following configurations:

- UTA/UTA2 cards ordered at the same time as the controller are configured before shipment to have the personality you requested.
- UTA/UTA2 cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard UTA/UTA2 ports on new controllers are configured (before shipment) to have the personality you requested.

However, you should check the configuration of the UTA/UTA2 ports on node4 and change it, if necessary.



Enter the commands in this section at the Maintenance mode shell prompt.

Steps

1. Check how the ports are currently configured on node4:

```
system node hardware unified-connect show
```

The system displays output similar to the following example:

```
*> ucadmin show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
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.

- Verify the settings:

```
ucadmin show
```

Examine the output of the `ucadmin show` command and determine whether the UTA/UTA2 ports have the personality you want.

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`:

```
*> ucadmin show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type
f-a	1a	fc	initiator	-	-
f-a	1b	fc	target	-	initiator
f-a	2a	fc	target	cna	-
f-a	2b	fc	target	cna	-

4 entries were displayed.

```
*>
```

- 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 9 and go to Step 10 .

5. Take one of the following actions:

If you are configuring...	Then...
Ports on a UTA/UTA2 card	Go to Step 6
Onboard UTA/UTA2 ports	Skip Step 6 and go to Step 7 .

6. If the adapter is in initiator mode, and if the UTA/UTA2 port is online, take the UTA/UTA2 port offline:

```
storage disable adapter adapter_name
```

Adapters in target mode are automatically offline in Maintenance mode.

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

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

- `-m` is the personality mode, FC or 10GbE UTA.
- `-t` is the FC4 type, target or initiator.



You must use FC initiator for tape drives and MetroCluster configurations. You must use the FC target for SAN clients.

8. Place any target ports online by entering the following command, once for each port:

```
storage enable adapter <adapter_name>
```

9. Cable the port.

10. Exit Maintenance mode:

```
halt
```

11. Boot the node into boot menu:

```
boot_ontap menu
```

What's next?

- If you are upgrading to an AFF A800 system, go to [Reassign node 2 disks to node 4, Step 9](#).
- For all other system upgrades, go to [Reassign node2 disks to node4, Step 1](#).

Reassign node2 disks to node4

You need to reassign the disks that belonged to node2 to node4 before verifying the node4 installation..

Steps

1. Verify that node2 has stopped at the boot menu and reassign the disks of node2 to node4:

```
boot_after_controller_replacement
```

After a short delay, you are prompted to enter the name of the node that is being replaced. If there are shared disks (also called Advanced Disk Partitioning (ADP) or partitioned disks), you are prompted to enter the node name of the HA partner.

These prompts might get buried in the console messages. If you do not enter a node name or enter an incorrect name, you are prompted to enter the name again.

Expand the console output example

```
LOADER-A> boot_ontap menu
.
.
<output truncated>
.
All rights reserved.
*****
*                                     *
* Press Ctrl-C for Boot Menu. *
*                                     *
*****
.
<output truncated>
.
Please choose one of the following:
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 22/7
(22/7)                                     Print this secret List
(25/6)                                     Force boot with multiple filesystem
disks missing.
(25/7)                                     Boot w/ disk labels forced to clean.
(29/7)                                     Bypass media errors.
(44/4a)                                    Zero disks if needed and create new
flexible root volume.
(44/7)                                     Assign all disks, Initialize all
disks as SPARE, write DDR labels
.
.
<output truncated>
.
.
(wipeconfig)                               Clean all configuration on boot
device
(boot_after_controller_replacement) Boot after controller upgrade
```

```

(boot_after_mcc_transition)      Boot after MCC transition
(9a)                             Unpartition all disks and remove
their ownership information.
(9b)                             Clean configuration and
initialize node with partitioned disks.
(9c)                             Clean configuration and
initialize node with whole disks.
(9d)                             Reboot the node.
(9e)                             Return to main boot menu.
The boot device has changed. System configuration information could
be lost. Use option (6) to
restore the system configuration, or option (4) to initialize all
disks and setup a new system.
Normal Boot is prohibited.
Please choose one of the following:
(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? boot_after_controller_replacement
This will replace all flash-based configuration with the last backup
to disks. Are you sure
you want to continue?: yes
.
.
<output truncated>
.
.
Controller Replacement: Provide name of the node you would like to
replace:
<nodename of the node being replaced>
Controller Replacement: Provide High Availability partner of node1:
<nodename of the partner of the node being replaced>
Changing sysid of node node2 disks.
Fetched sanown old_owner_sysid = 536940063 and calculated old sys id
= 536940063
Partner sysid = 4294967295, owner sysid = 536940063
.
.

```

```

<output truncated>
.
.
varfs_backup_restore: restore using /mroot/etc/varfs.tgz
varfs_backup_restore: attempting to restore /var/kmip to the boot
device
varfs_backup_restore: failed to restore /var/kmip to the boot device
varfs_backup_restore: attempting to restore env file to the boot
device
varfs_backup_restore: successfully restored env file to the boot
device wrote
    key file "/tmp/rndc.key"
varfs_backup_restore: timeout waiting for login
varfs_backup_restore: Rebooting to load the new varfs
Terminated
<node reboots>
System rebooting...
.
.
Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...
Rebooting to load the restored env file...
.
System rebooting...
.
.
.
<output truncated>
.
.
.
.
WARNING: System ID mismatch. This usually occurs when replacing a
boot device or NVRAM cards!
Override system ID? {y|n} y
.
.
.
.
Login:

```



In the above console output example, ONTAP will prompt you for the partner node name if the system uses Advanced Disk Partitioning (ADP) disks.

2. If the system goes into a reboot loop with the message `no disks found`, it indicates that the system has reset the FC or UTA/UTA2 ports back to the target mode and therefore is unable to see any disks. Select one of the following tasks to resolve this issue:
 - Perform [Step 3](#) to [Step 8](#) on node4
 - Go to section [Verify the node4 installation](#)
3. Press Ctrl-C during AUTOBOOT to stop the node at the LOADER> prompt.
4. At the LOADER prompt, enter maintenance mode:

```
boot_ontap maint
```

5. In maintenance mode, display all the previously set initiator ports that are now in target mode:

```
ucadmin show
```

Change the ports back to initiator mode:

```
ucadmin modify -m fc -t initiator -f adapter name
```

6. Verify that the ports have been changed to initiator mode:

```
ucadmin show
```

7. Exit maintenance mode:

```
halt
```



If you are upgrading from a system that supports external disks to a system that also supports external disks, go to [Step 8](#).

If you are upgrading from a system that uses external disks to a system that supports both internal and external disks, for example, an AFF A800 system, go to [Step 9](#).

8. At the LOADER prompt, boot up:

```
boot_ontap menu
```

Now, on booting, the node can detect all the disks that were previously assigned to it and can boot up as expected.

When the cluster nodes you are replacing use root volume encryption, ONTAP is unable to read the volume information from the disks. Restore the keys for the root volume.



This only applies when the root volume is using NetApp Volume Encryption.

- a. Return to the special boot menu:

```
LOADER> boot_ontap menu
```


Please choose one of the following:

- (1) Normal Boot.
- (2) Boot without /etc/rc.
- (3) Change password.
- (4) Clean configuration and initialize all disks.
- (5) Maintenance mode boot.
- (6) Update flash from backup config.
- (7) Install new software first.
- (8) Reboot node.
- (9) Configure Advanced Drive Partitioning.
- (10) Set Onboard Key Manager recovery secrets.
- (11) Configure node for external key management.

Selection (1-11)? 10

b. Select **(10) Set Onboard Key Manager recovery secrets**

c. Enter `y` at the following prompt:

```
This option must be used only in disaster recovery procedures. Are you sure?  
(y or n): y
```

d. At the prompt, enter the key-manager passphrase.

e. Enter the backup data when prompted.



You must have obtained the passphrase and backup data in the [Prepare the nodes for upgrade](#) section of this procedure.

f. After the system boots to the special boot menu again, run option **(1) Normal Boot**



You might encounter an error at this stage. If an error occurs, repeat the substeps in [Step 8](#) until the system boots normally.

9. If you are upgrading from a system with external disks to a system that supports internal and external disks (AFF A800 systems, for example), set the node2 aggregate as the root aggregate to ensure node4 boots from the root aggregate of node2. To set the root aggregate, go to the boot menu on node4 and select option 5 to enter maintenance mode.



You must perform the following substeps in the exact order shown; failure to do so might cause an outage or even data loss.

The following procedure sets node4 to boot from the root aggregate of node2:

a. Enter maintenance mode:

```
boot_ontap maint
```

b. Check the RAID, plex, and checksum information for the node2 aggregate:

```
aggr status -r
```

- c. Check the status of the node2 aggregate:

```
aggr status
```

- d. If necessary, bring the node2 aggregate online:

```
aggr_online root_aggr_from_node2
```

- e. Prevent the node4 from booting from its original root aggregate:

```
aggr offline root_aggr_on_node4
```

- f. Set the node2 root aggregate as the new root aggregate for node4:

```
aggr options aggr_from_node2 root
```

- g. Verify that the root aggregate of node4 is offline and the root aggregate for the disks brought over from node2 is online and set to root:

```
aggr status
```



Failing to perform the previous substep might cause node4 to boot from the internal root aggregate, or it might cause the system to assume a new cluster configuration exists or prompt you to identify one.

The following shows an example of the command output:

```
-----  
Aggr State                               Status                               Options  
aggr 0_nst_fas8080_15 online             raid_dp, aggr                      root, nosnap=on  
                                           fast zeroed  
                                           64-bit  
aggr0 offline                           raid_dp, aggr                      diskroot  
                                           fast zeroed`  
                                           64-bit  
-----
```

Verify the node4 installation

You must verify that the physical ports from node2 map correctly to the physical ports on node4. This will enable node4 to communicate with other nodes in the cluster and with the network after the upgrade.

About this task

Refer to [References](#) to link to the *Hardware Universe* to capture information about the ports on the new nodes. You will use the information later in this section.

Physical port layout might vary, depending on the model of the nodes. When the new node boots up, ONTAP will try to determine which ports should host cluster LIFs in order to automatically come into quorum.

If the physical ports on node2 do not map directly to the physical ports on node4, the subsequent section [Restore network configuration on node4](#) must be used to repair network connectivity.

After you install and boot node4, you must verify that it is installed correctly. You must wait for node4 to join quorum and then resume the relocation operation.

At this point in the procedure, the operation will have paused as node4 joins quorum.

Steps

- 1. Verify that node4 has joined quorum:

```
cluster show -node node4 -fields health
```

The output of the `health` field should be `true`.

- 2. Verify that node4 is part of the same cluster as node3 and that it is healthy:

```
cluster show
```

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

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

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

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

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

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

Repeat Step 3 to verify that the cluster LIF is now listening on port 7700.

5. Switch to advanced privilege mode:

```
set advanced
```

6. Check the status of the controller replacement operation and verify that it is in a paused state and in the same state it was in before node2 was halted to perform the physical tasks of installing new controllers and moving cables:

```
system controller replace show
```

```
system controller replace show-details
```

7. If you are working on a MetroCluster system, verify that the replaced controller is configured correctly for the MetroCluster configuration; the MetroCluster configuration should be in a healthy state. Refer to [Verify the health of the MetroCluster configuration](#).

Reconfigure the intercluster LIFs on MetroCluster node node4, and check cluster peering to restore communication between the MetroCluster nodes before proceeding to [Step 6](#).

Check the MetroCluster node status:

```
metrocluster node show
```

8. Resume the controller replacement operation:

```
system controller replace resume
```

9. Controller replacement will pause for intervention with the following message:

```

Cluster::*> system controller replace show
Node                Status                Error-Action
-----
Node2(now node4) Paused-for-intervention  Follow the instructions
given in
Step Details
Node2
Step Details:
-----
To complete the Network Reachability task, the ONTAP network
configuration must be
manually adjusted to match the new physical network configuration of the
hardware.
This includes:

1. Re-create the interface group, if needed, before restoring VLANs. For
detailed
commands and instructions, refer to the "Re-creating VLANs, ifgrps, and
broadcast
domains" section of the upgrade controller hardware guide for the ONTAP
version
running on the new controllers.
2. Run the command "cluster controller-replacement network displaced-
vlangs show"
to check if any VLAN is displaced.
3. If any VLAN is displaced, run the command "cluster controller-
replacement
network displaced-vlangs restore" to restore the VLAN on the desired
port.
2 entries were displayed.

```



In this procedure, section *Re-creating VLANs, ifgrps, and broadcast domains* has been renamed *Restoring network configuration on node4*.

10. With the controller replacement in a paused state, proceed to the next section of this document to restore network configuration on the node.

Restore network configuration on node4

After you confirm that node4 is in quorum and can communicate with node3, verify that node2's VLANs, interface groups and broadcast domains are seen on node4. Also, verify that all node4 network ports are configured in their correct broadcast domains.

About this task

For more information on creating and re-creating VLANs, interface groups, and broadcast domains, refer to [References](#) to link to *Network Management*.



If you are changing the port speed of the e0a and e1a cluster ports on AFF A800 or AFF C800 systems, you might observe malformed packets being received after the speed conversion. See [NetApp Bugs Online Bug ID 1570339](#) and the knowledge base article [CRC errors on T6 ports after converting from 40GbE to 100GbE](#) for guidance.

Steps

1. List all the physical ports that are on upgraded node2 (referred to as node4):

```
network port show -node node4
```

All physical network ports, VLAN ports and interface group ports on the node are displayed. From this output you can see any physical ports that have been moved into the `Cluster` broadcast domain by ONTAP. You can use this output to aid in deciding which ports should be used as interface group member ports, VLAN base ports or standalone physical ports for hosting LIFs.

2. List the broadcast domains on the cluster:

```
network port broadcast-domain show
```

3. List the network port reachability of all ports on node4:

```
network port reachability show
```

The output from the command looks similar to the following example:

```
clusterA::*> reachability show -node node2_node4
(network port reachability show)
```

Node	Port	Expected Reachability	Reachability Status

node2_node4			
	a0a	Default:Default	no-reachability
	a0a-822	Default:822	no-reachability
	a0a-823	Default:823	no-reachability
	e0M	Default:Mgmt	ok
	e0a	Cluster:Cluster	misconfigured-
reachability	e0b	Cluster:Cluster	no-reachability
	e0c	Cluster:Cluster	no-reachability
	e0d	Cluster:Cluster	no-reachability
	e0e	Cluster:Cluster	ok
	e0e-822	-	no-reachability
	e0e-823	-	no-reachability
	e0f	Default:Default	no-reachability
	e0f-822	Default:822	no-reachability
	e0f-823	Default:823	no-reachability
	e0g	Default:Default	misconfigured-
reachability	e0h	Default:Default	ok
	e0h-822	Default:822	ok
	e0h-823	Default:823	ok

18 entries were displayed.

In the above example, node2_node4 is just booted after controller replacement. It has several ports that have no reachability and are pending a reachability scan.

- Repair the reachability for each of the ports on node4 with a reachability status other than `ok`. Run the following command, first on any physical ports, then on any VLAN ports, one at a time:

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

The output looks like the following example:

```
Cluster ::> reachability repair -node node2_node4 -port e0h
```

```
Warning: Repairing port "node2_node4: e0h" may cause it to move into a
different broadcast domain, which can cause LIFs to be re-homed away
from the port. Are you sure you want to continue? {y|n}:
```

A warning message, as shown above, is expected for ports with a reachability status that might be different from the reachability status of the broadcast domain where it is currently located.

Review the connectivity of the port and answer *y* or *n* as appropriate.

Verify that all physical ports have their expected reachability:

```
network port reachability show
```

As the reachability repair is performed, ONTAP attempts to place the ports in the correct broadcast domains. However, if a port's reachability cannot be determined and does not belong to any of the existing broadcast domains, ONTAP will create new broadcast domains for these ports.

5. If interface group configuration does not match the new controller physical port layout, modify it by using the following steps.
 - a. You must first remove physical ports that should be interface group member ports from their broadcast domain membership. You can do this by using the following command:

```
network port broadcast-domain remove-ports -broadcast-domain  
broadcast_domain_name -ports node_name:port_name
```

- b. Add a member port to an interface group:

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

- c. The interface group is automatically added to the broadcast domain about a minute after the first member port is added.
 - d. Verify that the interface group was added to the appropriate broadcast domain:

```
network port reachability show -node node_name -port ifgrp
```

If the interface group's reachability status is not *ok*, assign it to the appropriate broadcast domain:

```
network port broadcast-domain add-ports -broadcast-domain  
broadcast_domain_name -ports node:port
```

6. Assign appropriate physical ports to the *Cluster* broadcast domain:

- a. Determine which ports have reachability to the *Cluster* broadcast domain:

```
network port reachability show -reachable-broadcast-domains Cluster:Cluster
```

- b. Repair any port with reachability to the *Cluster* broadcast domain, if its reachability status is not *ok*:

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

7. Move the remaining physical ports into their correct broadcast domains by using one of the following commands:

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

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



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

Verify that there are no unreachable or unexpected ports present. Check the reachability status for all physical ports by using the following command and examining the output to confirm the status is `ok`:

```
network port reachability show -detail
```

8. Restore any VLANs that might have become displaced by using the following steps:

a. List displaced VLANs:

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

Output like the following should display:

```
Cluster::*> displaced-vlans show
(cluster controller-replacement network displaced-vlans show)

      Original
Node      Base Port      VLANs
-----
Node1     a0a            822, 823
          e0e            822, 823
```

b. Restore VLANs that were displaced from their previous base ports:

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

The following is an example of restoring VLANs that have been displaced from interface group `a0a` back onto the same interface group:

```
Cluster::*> displaced-vlans restore -node node2_node4 -port a0a
-destination-port a0a
```

The following is an example of restoring displaced VLANs on port `"e0e"` to `"e0h"`:

```
Cluster::*> displaced-vlans restore -node node2_node4 -port e0e
-destination-port e0h
```

When a VLAN restore is successful, the displaced VLANs are created on the specified destination port. The VLAN restore fails if the destination port is a member of an interface group, or if the destination port is down.

Wait about one minute for newly restored VLANs to be placed into their appropriate broadcast domains.

c. Create new VLAN ports as needed for VLAN ports that are not in the `cluster controller-replacement network displaced-vlans show` output but should be configured on other

physical ports.

9. Delete any empty broadcast domains after all port repairs have been completed:

```
network port broadcast-domain delete -broadcast-domain broadcast_domain_name
```

10. Verify port reachability:

```
network port reachability show
```

When all ports are correctly configured and added to the correct broadcast domains, the `network port reachability show` command should report the reachability status as `ok` for all connected ports, and the status as `no-reachability` for ports with no physical connectivity. If any ports report a status other than these two, perform the reachability repair and add or remove ports from their broadcast domains as instructed in [Step 4](#).

11. Verify that all ports have been placed into broadcast domains:

```
network port show
```

12. Verify that all ports in the broadcast domains have the correct maximum transmission unit (MTU) configured:

```
network port broadcast-domain show
```

13. Restore LIF home ports, specifying the Vserver(s) and LIF(s) home ports, if any, that need to be restored:

- a. List any LIFs that are displaced:

```
displaced-interface show
```

- b. Restore LIF home ports:

```
displaced-interface restore-home-node -node node_name -vserver vserver_name  
-lif-name LIF_name
```

14. Verify that all LIFs have a home port and are administratively up:

```
network interface show -fields home-port, status-admin
```

Restore key-manager configuration on node4

If you are using NetApp Volume Encryption (NVE) and NetApp Aggregate Encryption (NAE) to encrypt volumes on the system you are upgrading, the encryption configuration must be synchronized to the new nodes. If you do not synchronize the key-manager, when you relocate the node2 aggregates from node3 to node4 by using ARL, failures might occur because node4 does not have the required encryption keys to bring encrypted volumes and aggregates online.

About this task

Synchronize the encryption configuration to the new nodes by performing the following steps:

Steps

1. Run the following command from node4:

```
security key-manager onboard sync
```

2. Verify that the SVM-KEK key is restored to "true" on node4 before you relocate the data aggregates:

```
::> security key-manager key query -node node4 -fields restored -key  
-type SVM-KEK
```

Example

```
::> security key-manager key query -node node4 -fields restored -key  
-type SVM-KEK
```

node	vserver	key-server	key-id
restored			
-----	-----	-----	-----
node4	svm1	""	0000000000000000020000000000a008a81976
true			2190178f9350e071fbb90f00000000000000000

Move non-root aggregates and NAS data LIFs owned by node2 from node3 to node4

After verifying the network configuration on node4, you need to relocate the NAS data LIFs owned by node2 from node3 to node4 and confirm that the SAN LIFs exist 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 the correct ports after you bring node4 online.



If you are changing the port speed of the T6-based Ethernet network interface cards or motherboard ports, you might observe malformed packets being received after the speed conversion. See [NetApp Bugs Online Bug ID 1570339](#) and the knowledge base article [CRC errors on T6 ports after converting from 40GbE to 100GbE](#) for guidance.

Steps

1. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

- Cluster quorum check
- System ID check
- Image version check
- Target platform check
- Network reachability check

The system pauses the operation at this stage in the network reachability check.

2. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following checks:

- Cluster health check
- Cluster LIF status check

After performing these checks, the system relocates the non-root aggregates and NAS data LIFs owned by node2 to the new controller, node4.

The controller replacement operation pauses after the resource relocation is complete.

3. Check the status of the aggregate relocation and NAS data LIF move operations:

```
system controller replace show-details
```

If the controller replacement procedure is paused, check and correct the error, if any, and then issue `resume` to continue the operation.

4. If necessary, restore and revert displaced LIFs or manually migrate and modify the node2 LIFs that failed to relocate automatically to node4.

Restore and revert displaced LIFs

- a. List any displaced LIFs:

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

- b. If any LIFs are displaced, restore the home node back to node4:

```
cluster controller-replacement network displaced-interface  
restore-home-node -node <node4_nodename> -vserver <vserver name>  
-lif-name <lif_name>
```

Manually migrate and modify LIFs

- a. Migrate the LIFs that failed to relocate automatically to node4:

```
network interface migrate -vserver <vserver name> -lif <lif_name>  
-destination-node <node4_nodename> -destination-port  
<port_on_node4>
```

- b. Modify the home node and home port for the migrated LIFs:

```
network interface modify -vserver <vserver_name> -lif  
<data_lif_name> -home-node <node4_nodename> -home-port  
<home_port>
```

5. Resume the operation to prompt the system to perform the required post-checks:

```
system controller replace resume
```

The system performs the following post-checks:

- Cluster quorum check
- Cluster health check
- Aggregates reconstruction check
- Aggregate status check
- Disk status check
- Cluster LIF status check
- Volume check

Stage 6. Complete the upgrade

Manage authentication using KMIP servers

With ONTAP 9.8 or later, you can use Key Management Interoperability Protocol (KMIP) servers to manage authentication keys.

Steps

1. Add a new controller:

```
security key-manager external enable
```

2. Add the key manager:

```
security key-manager external add-servers -key-servers  
key_management_server_ip_address
```

3. Verify that the key management servers are configured and available to all nodes in the cluster:

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

4. Restore the authentication keys from all linked key management servers to the new node:

```
security key-manager external restore -node new_controller_name
```

Confirm that the new controllers are set up correctly

To confirm 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 confirm that node3 owns node1's aggregates and that node4 owns node2's aggregates, and that the volumes for both nodes are online.

Steps

1. After the post-checks of node2, the storage failover and cluster HA pair for the node2 cluster are enabled. When the operation is done, both nodes show as completed and the system performs some cleanup operations.
2. Verify that storage failover is enabled:

```
storage failover show
```

The following example shows the output of the command when storage failover is enabled:

```
cluster::> storage failover show
```

Takeover			
Node	Partner	Possible	State Description
-----	-----	-----	-----
node3	node4	true	Connected to node4
node4	node3	true	Connected to node3

3. Verify that node3 and node4 belong to the same cluster by using the following command and examining the output:

```
cluster show
```

4. Verify that node3 and node4 can access each other's storage by using the following command and examining the output:

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

5. Verify that neither node3 nor node4 owns data LIFs home-owned by other nodes in the cluster by using the following command and examining the output:

```
network interface show
```

If neither node3 or node4 owns data LIFs home-owned by other nodes in the cluster, revert the data LIFs to their home owner:

```
network interface revert
```

6. Verify that node3 owns the aggregates from node1 and that node4 owns the aggregates from node2:

```
storage aggregate show -owner-name <node3>
```

```
storage aggregate show -owner-name <node4>
```

7. Determine whether any volumes are offline:

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

```
volume show -node <node4> -state offline
```

8. If any volumes are offline, compare them with the list of offline volumes that you captured in the section [Prepare the nodes for upgrade](#), and bring online any of the offline volumes, as required, by using the following command, once for each volume:

```
volume online -vserver <vserver_name> -volume <volume_name>
```

9. Install new licenses for the new nodes by using the following command 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 multiple licenses at once, separating each license key by a comma.

10. Remove all of the old licenses from the original nodes by using one of the following commands:

```
system license clean-up -unused -expired
```

```
system license delete -serial-number <node_serial_number> -package  
<licensable_package>
```

- Delete all expired licenses:

```
system license clean-up -expired
```

- Delete all unused licenses:

```
system license clean-up -unused
```

- Delete a specific license from a cluster by using 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.

11. Verify that the licenses are correctly installed by using the following command and examining the output:

```
system license show
```

You can compare the output with the output that you captured in the section [Prepare the nodes for upgrade](#).

12. If self-encrypting drives are being used in the configuration and you have set the `kmip.init.maxwait` variable to `off` (for example, in [Install and boot node4, Step 22](#)), you must unset the variable:

```
set diag; systemshell -node node_name -command sudo kenv -u -p  
kmip.init.maxwait
```

13. Configure the SPs by using the following command on both nodes:

```
system service-processor network modify -node node_name
```

Refer to [References](#) to link to the *System Administration Reference* for information about the SPs and the *ONTAP 9.8 Commands: Manual Page Reference* for detailed information about the `system service-processor network modify` command.

14. If you want to set up a switchless cluster on the new nodes, refer to [References](#) to link to the *NetApp Support Site* and follow the instructions in *Transitioning to a two-node switchless cluster*.

After you finish

If Storage Encryption is enabled on node3 and node4, complete the section [Set up Storage Encryption on the new controller module](#). Otherwise, complete the section [Decommission the old system](#).

Set 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. You must 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 external show-status
```

```
security key-manager onboard show-backup
```

2. Add the key management servers listed in the previous step to the key management server list in the new controller.

- a. Add the key management server:

```
security key-manager external add-servers -key-servers  
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 external 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 external enable
```

- 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 external restore -node new_controller_name
```

Set up NetApp Volume or Aggregate Encryption on the new controller module

If the replaced controller or high availability (HA) partner of the new controller uses NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE), you must configure the new controller module for NVE or NAE.

About this task

This procedure includes steps that are performed on the new controller module. You must enter the command on the correct node.

Onboard Key Manager

Configure NVE or NAE using the Onboard Key Manager.

Steps

1. Restore authentication keys from all linked key management servers to the new node:

```
security key-manager onboard sync
```

External Key Management

Configure NVE or NAE using External Key Management.

Steps

1. Verify that the key management servers are still available, their status, and their authentication key information:

```
security key-manager key query -node node
```

2. Add the key management servers listed in the previous step to the key management server list in the new controller:

- a. Add the key management server:

```
security key-manager external add-servers -key-servers  
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 external 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 external enable
```

- 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 external restore
```

This command needs the OKM passphrase

For more information, see the Knowledge Base article [How to restore external key manager server configuration from the ONTAP boot menu](#).

After you finish

Check if any volumes were taken offline because authentication keys were not available or EKM servers could not be reached. Bring those volumes back online by using the `volume online` command.

Decommission 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. Refer to [References](#) to link to the *NetApp Support Site* and log in.
2. Select **Products > My Products** from the menu.
3. On the **View Installed Systems** page, choose which **Selection Criteria** you want to use to display information about your system.

You can choose one of the following to locate your system:

- Serial Number (located on the back of the unit)
- Serial Numbers for My Location

4. Select **Go!**

A table displays cluster information, including the serial numbers.

5. Locate the cluster in the table and select **Decommission this system** from the Product Tool Set drop-down menu.

Resume SnapMirror operations

You can resume SnapMirror transfers that were quiesced before upgrade and resume the SnapMirror relationships. The updates are on schedule after the upgrade is completed.

Steps

1. Verify the SnapMirror status on the destination:

```
snapmirror show
```

2. Resume the SnapMirror relationship:

```
snapmirror resume -destination-vserver vservers_name
```

Troubleshoot

Aggregate relocation failures

Aggregate relocation (ARL) might fail at different points during the upgrade.

Check for aggregate relocation failure

During the procedure, ARL might fail in Stage 2, Stage 3, or Stage 5.

Steps

1. Enter the following command and examine the output:

```
storage aggregate relocation show
```

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, refer to [References](#) to link to the *ONTAP 9.8 Commands: Manual Page Reference*.

Aggregates originally on node1 are owned by node4 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 correctly, 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 [Prepare the nodes for 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 [Prepare the nodes for upgrade](#) and note any aggregates that were not correctly relocated.
3. Relocate the aggregates left behind on node4:

```
storage aggregate relocation start -node node4 -aggr aggr_node_1 -destination node3
```

Do not use the `-ndo-controller-upgrade` parameter during this relocation.

4. 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 the pre-check phase

Node1 or node2 crashes before the pre-check phase with HA pair still enabled

If either node1 or node2 crashes before the pre-check phase, no aggregates have been relocated yet and the HA pair 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.

Reboots, panics, or power cycles during first resource-release phase

Node1 crashes during the first resource-release phase with HA pair still enabled

Some or all aggregates have been relocated from node1 to node2, and HA pair is still enabled. Node2 takes over node1's root volume and any non-root aggregates that were not relocated.

About this task

Ownership of aggregates that were relocated look the same as the ownership of non-root aggregates that were taken over because the home owner has not changed.

When node1 enters the `waiting for giveback` state, node2 gives back all of the node1 non- root

aggregates.

Steps

1. After node1 is booted up, all the non-root aggregates of node1 have moved back to node1. You must perform a manual aggregate relocation of the aggregates from node1 to node2:

```
storage aggregate relocation start -node node1 -destination node2 -aggregate  
-list * -ndocontroller-upgrade true
```
2. Continue with the node-pair upgrade procedure.

Node1 crashes during the first resource-release phase while HA pair is disabled

Node2 does 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.

Node2 fails during the first resource-release phase with HA pair still enabled

Node1 has relocated some or all of its aggregates to node2. The HA pair is enabled.

About this task

Node1 takes over all of node2's aggregates as well as any of its own aggregates that it had relocated to node2. When node2 boots up, the aggregate relocation is completed automatically.

Steps

1. Bring up node2.
2. Continue with the node-pair upgrade procedure.

Node2 crashes during the first resource-release phase and after HA pair is disabled

Node1 does not take over.

Steps

1. Bring up node2.

A client outage occurs for all aggregates while node2 is booting up.
2. Continue with the rest of the node-pair upgrade procedure.

Reboots, panics, or power cycles during the first verification phase

Node2 crashes during the first verification phase with HA pair disabled

Node3 does not take over following a node2 crash as the HA pair is already disabled.

Steps

1. Bring up node2.

A client outage occurs for all aggregates while node2 is booting up.
2. Continue with the node-pair upgrade procedure.

Node3 crashes during the first verification phase with HA pair disabled

Node2 does 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.

Reboots, panics, or power cycles during first resource-regain phase

Node2 crashes during the first resource-regain phase during aggregate relocation

Node2 has relocated some or all of its aggregates from node1 to node3. Node3 serves data from aggregates that were relocated. The HA pair is disabled and hence there is no takeover.

About this task

There is client outage for aggregates that were not relocated. On booting up node2, the aggregates of node1 are relocated to node3.

Steps

1. Bring up node2.
2. Continue with the node-pair upgrade procedure.

Node3 crashes during the first resource-regain phase during aggregate relocation

If node3 crashes while node2 is relocating aggregates to node3, the task continues after node3 boots up.

About this task

Node2 continues to serve remaining aggregates, but aggregates that were already relocated to node3 encounter client outage while node3 is booting up.

Steps

1. Bring up node3.
2. Continue with the controller upgrade.

Reboots, panics, or power cycles during post-check phase

Node2 or node3 crashes during the post-check phase

The HA pair is disabled hence this is no takeover. There is a client outage for aggregates belonging to the node that rebooted.

Steps

1. Bring up the node.
2. Continue with the node-pair upgrade procedure.

Reboots, panics, or power cycles during second resource-release phase

Node3 crashes during the second resource-release phase

If node3 crashes while node2 is relocating aggregates, the task continues after node3 boots up.

About this task

Node2 continues to serve remaining aggregates but aggregates that were already relocated to node3 and node3's own aggregates encounter client outages while node3 is booting.

Steps

1. Bring up node3.
2. Continue with the controller upgrade procedure.

Node2 crashes during the second resource-release phase

If node2 crashes during aggregate relocation, node2 is not taken over.

About this task

Node3 continues to serve the aggregates that have been relocated, but the aggregates owned by node2 encounter client outages.

Steps

1. Bring up node2.
2. Continue with the controller upgrade procedure.

Reboots, panics, or power cycles during the second verification phase**Node3 crashes during the second verification phase**

If node3 crashes during this phase, takeover does not happen because the HA pair is already disabled.

About this task

There is a client outage for all aggregates until node3 reboots.

Steps

1. Bring up node3.
2. Continue with the node-pair upgrade procedure.

Node4 crashes during the second verification phase

If node4 crashes during this phase, takeover does not happen. Node3 serves data from the aggregates.

About this task

There is an outage for non-root aggregates that were already relocated until node4 reboots.

Steps

1. Bring up node4.
2. Continue with the node-pair upgrade procedure.

Issues that can arise in multiple stages of the procedure

Some issues can occur during 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	
		Possible	State Description
node1	node2	false	Unknown
node2	node1	false	Node owns partner aggregates as part of the non-disruptive head upgrade procedure. Takeover is not possible: Storage failover is disabled.

The following example shows the output of the `storage failover show` command after a reboot or panic:

```
cluster::> storage failover show
```

Node	Partner	Takeover	
		Possible	State Description
node1	node2	-	Unknown
node2	node1	false	Waiting for node1, Partial giveback, Takeover is not possible: Storage failover is disabled

Although the output says that a node is in partial giveback and that storage failover is disabled, you can disregard this message.

Steps

No action is required; continue with the node-pair upgrade procedure.

LIF migration failure

After you migrate LIFs, they might not come online after migration in Stage 2, Stage 3, or Stage 5.

Steps

- 1. Verify that the 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.

2. Check the physical connectivity of the network cable if the physical state of the port is down.

References

When performing the procedures in this content, you might need to consult reference content or go to reference websites.

Reference content

Content specific to this upgrade are listed in the table below.

Content	Description
Administration overview with the CLI	Describes how to administer ONTAP systems, shows you how to use the CLI interface, how to access the cluster, how to manage nodes, and much more.
Decide whether to use System Manager or the ONTAP CLI for cluster setup	Describes how to set up and configure ONTAP.
Disk and aggregate management with the CLI	Describes how to manage ONTAP physical storage using the CLI. It shows you how to create, expand, and manage aggregates, how to work with Flash Pool aggregates, how to manage disks, and how to manage RAID policies.
HA pair management	Describes how to install and manage high-availability clustered configurations, including storage failover and takeover/giveback.
Logical storage management with the CLI	Describes how to efficiently manage your logical storage resources, using volumes, FlexClone volumes, files, and LUNs, FlexCache volumes, deduplication, compression, qtrees, and quotas.
MetroCluster Management and Disaster Recovery	Describes how to perform MetroCluster switchover and switchback operations, both in planned maintenance operations, or in the event of a disaster.
MetroCluster Upgrade and Expansion	Provides procedures for upgrading controller and storage models in the MetroCluster configuration, transitioning from a MetroCluster FC to a MetroCluster IP configuration, and expanding the MetroCluster configuration by adding additional nodes.
Network Management	Describes how to configure and manage physical and virtual network ports (VLANs and interface groups), LIFs, routing, and host-resolution services in clusters; optimize network traffic by load balancing; and monitor the cluster by using SNMP.
ONTAP 9.0 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.0 commands.
ONTAP 9.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.1 commands.

Content	Description
ONTAP 9.2 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.2 commands.
ONTAP 9.3 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.3 commands.
ONTAP 9.4 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.4 commands.
ONTAP 9.5 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.5 commands.
ONTAP 9.6 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.6 commands.
ONTAP 9.7 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.7 commands.
ONTAP 9.8 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.8 commands.
ONTAP 9.9.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.9.1 commands.
ONTAP 9.10.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.10.1 commands.
SAN management with the CLI	Describes how to configure and manage LUNs, igroups, and targets using the iSCSI and FC protocols, and namespaces and subsystems using the NVMe/FC protocol.
SAN configuration reference	Contains information about FC and iSCSI topologies and wiring schemes.
Upgrade by moving volumes or storage	Describes how to quickly upgrade controller hardware in a cluster by moving storage or volumes. Also describes how to convert a supported model to a disk shelf.
Upgrade ONTAP	Contains instructions for downloading and upgrading ONTAP.
Use "system controller replace" commands to upgrade controller hardware introduced in ONTAP 9.15.1 and later	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers introduced in ONTAP 9.15.1 and later by using "system controller replace" commands.
Use "system controller replace" commands to upgrade controller models in the same chassis	Describes the aggregate relocation procedures needed to non-disruptively upgrade a system, keeping the old system chassis and disks.
Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers running ONTAP 9.8 by using "system controller replace" commands.
Use aggregate relocation to manually upgrade controller hardware running ONTAP 9.8 or later	Describes the aggregate relocation procedures needed to perform manual non-disruptive controller upgrades running ONTAP 9.8 or later.
Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.5 to ONTAP 9.7	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers running ONTAP 9.5 to ONTAP 9.7 by using "system controller replace" commands.

Content	Description
Use aggregate relocation to manually upgrade controller hardware running ONTAP 9.7 or earlier	Describes the aggregate relocation procedures needed to perform manual non-disruptive controller upgrades running ONTAP 9.7 or earlier.

Reference sites

The [NetApp Support Site](#) also contains documentation about network interface cards (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.

Access [ONTAP 9 documentation](#).

Access the [Active IQ Config Advisor](#) tool.

Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.5 to 9.7

Learn about this ARL upgrade procedure

There are several aggregate relocation (ARL) methods for upgrading controller hardware. This procedure describes how to upgrade the controller hardware using ARL with "system controller replace commands" on systems running ONTAP 9.7, 9.6, or 9.5.

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 confirm 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.

Terminology used in this information

In this information, the original nodes are called "node1" and "node2", and the new nodes are called "node3" and "node4". During the described procedure, "node1" is replaced by "node3", and "node2" is replaced by "node4".

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

Important information:

- This procedure is complex and assumes that you have advanced ONTAP administration skills. You also must read and understand the [guidelines for upgrading controllers with ARL](#) and the [ARL upgrade sequence](#) before beginning the upgrade.
- 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.
- 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 to FAS systems and AFF systems.
- Beginning with ONTAP 9.6, this procedure applies to systems running 4-node MetroCluster configuration or higher. Because MetroCluster configuration sites can be at two physically different locations, the automated controller upgrade must be carried out individually at each MetroCluster site for an HA pair.
- If you are upgrading from an AFF A320 system, you can use volume moves to upgrade controller hardware or contact technical support. If you are willing to do volume moves, refer to [References](#) to link to *Upgrade by moving volumes or storage*.

Automate the controller upgrade process

During a controller upgrade, the controller is replaced with another controller running a newer or more powerful platform.

Earlier versions of this content contained instructions for a nondisruptive controller update process that was comprised of entirely manual steps. This content provides the steps for the new automated procedure.

The manual process was lengthy and complex but in this simplified procedure you can implement a controller update using aggregate relocation, which enables more efficient nondisruptive upgrades for HA pairs. There are significantly fewer manual steps, especially around validation, collection of information, and post checks.

Decide whether to use this aggregate relocation procedure

There are several aggregate relocation (ARL) methods for upgrading controller hardware. This procedure describes how to upgrade the controller hardware using ARL with "system controller replace commands" on systems running ONTAP 9.7, 9.6, or 9.5. You should only use this complex procedure if you're an experienced ONTAP administrator.

To help you decide if this ARL procedure is suitable for your controller hardware upgrade, you should review all of the following circumstances for supported upgrades:

- You are upgrading NetApp controllers running ONTAP 9.5, 9.6 or 9.7. This document is not applicable to upgrades to ONTAP 9.8.
- You don't 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.
- Your hardware upgrade combination is listed in the [supported model matrix](#).
- If you are upgrading a MetroCluster configuration, it is a 4-node or higher FC configuration, and all nodes are running ONTAP 9.6 or 9.7.



- If you're upgrading a system by swapping controller modules in the same chassis, such as AFF A800 or AFF C800, NetApp strongly recommends using the upgrade procedure that [upgrades controller models using ARL, keeping the existing system chassis and disks](#). This ARL procedure includes the steps that ensure the internal disks remain secure in the chassis when you remove and install the controllers during the upgrade procedure.

[Learn about the supported system upgrade combinations using ARL, keeping the existing system chassis and disks.](#)

- You can use NetApp Storage Encryption (NSE), NetApp Volume Encryption (NVE), and NetApp Aggregate Encryption (NAE) with this procedure.

Supported system upgrade combinations

The following tables shows the supported model matrix for the controller upgrade using this ARL procedure.

Old controller	Replacement controller
FAS8020, FAS8040, FAS8060, FAS8080	FAS8200, FAS8300, FAS8700, FAS9000
AFF8020, AFF8040, AFF8060, AFF8080	AFF A300, AFF A400, AFF A700 ¹ , AFF A800 ²
FAS8200	FAS8700, FAS9000, FAS8300 ^{4, 5}
AFF A300	AFF A700 ¹ , AFF A800 ^{2, 3} , AFF A400 ^{4, 5}



If your controller upgrade model combination is not in the above table, contact technical support.

¹ARL automated upgrade for the AFF A700 system is supported from ONTAP 9.7P2.

²If you are updating to an AFF A800 or a system that supports internal and external disks, you must follow specific instructions for the root aggregate on internal NVMe disks. Refer to [Reassign node1 disks to node3, Step 9](#) and [Reassign node2 disks to node4, Step 9](#).

³ARL automated upgrade from an AFF A300 to an AFF A800 system is supported from ONTAP 9.7P5.

⁴ARL automated upgrade from an AFF A300 to an AFF A400 and an FAS8200 to an FAS8300 system is supported from ONTAP 9.7P8.

⁵If you are upgrading from an AFF A300 to an AFF A400 or an FAS8200 to an FAS8300 system in a two-node switchless cluster configuration, you must pick temporary cluster ports for the controller upgrade. The AFF A400 and FAS8300 systems come in two configurations, as an Ethernet bundle where the mezzanine card ports are Ethernet type and as an FC bundle where the mezzanine ports are FC type.

- For an AFF A400 or an FAS8300 with an Ethernet type configuration, you can use any of the two mezzanine ports as temporary cluster ports.
- For an AFF A400 or an FAS8300 with an FC type configuration, you must add a four-port 10GbE network interface card (part number X1147A) to provide temporary cluster ports.
- After you complete a controller upgrade by using temporary cluster ports, you can nondisruptively migrate cluster LIFs to e3a and e3b, 100GbE ports on an AFF A400 system, and e0c and e0d, 100GbE ports on an FAS8300 system.

Choose a different hardware upgrade procedure

- [Review the alternative ARL methods available for upgrading controller hardware.](#)
- If you prefer a different method of upgrading the controller hardware and are willing to do volume moves, refer to [References](#) to link to *Upgrade by moving volumes or storage*.

Related information

Refer to [References](#) to link to the *ONTAP 9 Documentation*.

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 need the following tools to perform the upgrade:

- Grounding strap
- #2 Phillips screwdriver

Go to the [References](#) section to access the list of reference documents and reference sites required for this upgrade

Guidelines for upgrading controllers with ARL

To understand whether you can use aggregate relocation (ARL) to upgrade a pair of controllers running ONTAP 9.5 to ONTAP 9.7 depends on the platform and the configuration of both the original and replacement controllers.

Supported upgrades for ARL

When you upgrade a pair of nodes using this ARL procedure for ONTAP 9.5 to ONTAP 9.7, you must verify that 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. You must then compare the aggregate sizes and number of disks supported to the aggregate size and number of disks supported by the new system. Refer to [References](#) to link to the *Hardware Universe* where this information is available. 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. For more information about cluster mixing rules, refer to [References](#) to link to the *Hardware Universe*.



Before performing an AFF system upgrade, you must upgrade ONTAP to release versions 9.5P1 or later. These release levels are required for a successful upgrade.



If you are upgrading a system that supports internal drives (for example, an FAS2700 or AFF A250) but does NOT have internal drives, refer to [References](#) and use the procedure in the *Aggregate Relocation to Manually Upgrade Controller Hardware* content that is correct for your version of ONTAP.

If you are using ONTAP 9.6P11, 9.7P8, or later releases, it is recommended to enable Connectivity, Liveliness, and Availability Monitor (CLAM) takeover to return the cluster into quorum when certain node failures occur. The `kernel-service` command requires advanced privilege level access. For more information, see: [NetApp KB Article SU436: CLAM takeover default configuration changed](#).

Controller upgrade using ARL is supported on systems configured with SnapLock Enterprise and SnapLock Compliance volumes.

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.

Upgrades not supported for ARL

You cannot perform the following upgrades:

- To replacement controllers that do not support the disk shelves connected to the original controllers

Refer to [References](#) to link to the *Hardware Universe* for disk-support information.

- To entry level controllers with internal drives, for example: an FAS 2500.

If you want to upgrade entry level controllers with internal drives, refer to [References](#) to link to *Upgrade by moving volumes or storage* and go to the procedure *Upgrading a pair of nodes running clustered Data ONTAP by moving volumes*.

MetroCluster FC configuration

In a MetroCluster FC configuration, you must replace the disaster recovery/failover site nodes as soon as possible. Mismatch in controller models within a MetroCluster configuration isn't supported because controller model mismatch can cause disaster recovery mirroring to go offline. Use the `-skip-metrocluster-check true` command to bypass MetroCluster checks when you're replacing nodes at the second site.

Troubleshoot

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 table describing the different phases of the procedure in the section Overview of the ARL upgrade. Information about the failures that can occur is listed by the phase of the procedure.

If any problems occur while upgrading the controllers, refer to the [Troubleshoot](#) section. The information about failures that can occur is listed by the phase of the procedure in the [ARL upgrade sequence](#).

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

Verify the health of the MetroCluster configuration

Before starting an upgrade on a Fabric MetroCluster configuration, you must check the health of the MetroCluster configuration to verify proper operation.

Steps

1. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
dpgga-mcc-funct-8040-0403_siteA::*> metrocluster check run
```

The operation runs in the background.

2. After the `metrocluster check run` operation completes, view the results:

```
metrocluster check show
```

After approximately five minutes, the following results are displayed:

```
metrocluster_siteA::*> metrocluster check show
Last Checked On: 4/7/2019 21:15:05
Component          Result
-----
nodes               ok
lifs                ok
config-replication ok
aggregates          warning
clusters            ok
connections         not-applicable
volumes             ok
7 entries were displayed.
```

3. Check the status of the running MetroCluster check operation:

```
metrocluster operation history show -job-id 38
```

4. Verify that there are no health alerts:

```
system health alert show
```

Check for MetroCluster configuration errors

You can use the Active IQ Config Advisor tool available from the NetApp Support Site to check for common configuration errors.

If you do not have a MetroCluster configuration, you can skip this section.

About this task

Active IQ Config Advisor is a configuration validation and health check tool. You can deploy it at both secure sites and non-secure sites for data collection and system analysis.



Support for Config Advisor is limited, and available only online.

1. Download the [Active IQ Config Advisor](#) tool.
2. Run Active IQ Config Advisor, reviewing the output and following its recommendations to address any issues.

Verify switchover, healing, and switchback

You should verify the switchover, healing, and switchback operations of the MetroCluster configuration.

Refer to [References](#) to link to the *MetroCluster Management and Disaster Recovery* content and use the procedures mentioned for negotiated switchover, healing, and switchback.

Learn about the ARL upgrade sequence

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

Upgrade 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.

ARL upgrade sequence overview

During the procedure, you upgrade the original controller hardware with the replacement controller hardware, one controller at a time, taking advantage of the HA pair 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 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:

Stage	Steps
Stage 1. Prepare for the upgrade	<p>During Stage 1, you run prechecks and, if required, correct aggregate ownership. You must record certain information if you are managing storage encryption by using the Onboard Key Manager and you can choose to quiesce the SnapMirror relationships.</p> <p>Aggregate ownership at the end of Stage 1:</p> <ul style="list-style-type: none">• Node1 is the home owner and current owner of the node1 aggregates.• Node2 is the home owner and current owner of the node2 aggregates.

Stage	Steps
<p>Stage 2. Relocate and retire node1</p>	<p>During Stage 2, you relocate node1 non-root aggregates and NAS data LIFs to node2. This process is largely automated; the operation pauses to enable you to check its status. You must manually resume the operation. If required, you relocate failed or vetoed aggregates. You must record the necessary node1 information for use later in the procedure and then retire node1. You can also prepare to netboot node3 and node4 later in the procedure.</p> <p>Aggregate ownership at the end of Stage 2:</p> <ul style="list-style-type: none"> • Node2 is the current owner of node1 aggregates. • Node2 is the home owner and current owner of node2 aggregates.
<p>Stage 3. Install and boot node3</p>	<p>During Stage 3, you install and boot node3, map the cluster and node-management ports from node1 to node3, reassign the node1 disks to node3, and verify the node3 installation. If required, you set the FC or UTA/UTA2 configuration on node3 and confirm that node3 has joined quorum. You also relocate the node1 NAS data LIFs and non-root aggregates from node2 to node3 and verify that the SAN LIFs exist on node3.</p> <p>Aggregate ownership at the end of Stage 3:</p> <ul style="list-style-type: none"> • Node3 is the home owner and current owner of node1 aggregates. • Node2 is the home owner and current owner of node2 aggregates.
<p>Stage 4. Relocate and retire node2</p>	<p>During Stage 4, you relocate node2 non-root aggregates and non-SAN data LIFs to node3. You also record the necessary node2 information and then retire node2.</p> <p>Aggregate ownership at the end of Stage 4:</p> <ul style="list-style-type: none"> • Node3 is the home owner and current owner of aggregates that originally belonged to node1. • Node2 is the home owner of node2 aggregates. • Node3 is the current owner of node2 aggregates.
<p>Stage 5. Install and boot node4</p>	<p>During Stage 5, you install and boot node4, map the cluster and node-management ports from node2 to node4, reassign the node2 disks to node4, and verify the node4 installation. If required, you set the FC or UTA/UTA2 configuration on node4 and confirm that node4 has joined quorum. You also relocate node2 NAS data LIFs and non-root aggregates from node3 to node4 and verify the SAN LIFs exist on node4.</p> <p>Aggregate ownership at the end of Stage 5:</p> <ul style="list-style-type: none"> • Node3 is the home owner and current owner of the aggregates that originally belonged to node1. • Node4 is the home owner and current owner of aggregates that originally belonged to node2.

Stage	Steps
Stage 6. Complete the upgrade	During Stage 6, you confirm that the new nodes are set up correctly and, if the new nodes are encryption-enabled, you configure and set up Storage Encryption or NetApp Volume Encryption. You should also decommission the old nodes and resume the SnapMirror operations.

Stage 1. Prepare for upgrade

Prepare the nodes for upgrade

The controller replacement process begins with a series of prechecks. You also gather information about the original nodes for use later in the procedure and, if required, determine the type of self-encrypting drives that are in use.

Steps

1. Begin the controller replacement process by entering the following command in the ONTAP command line:

```
system controller replace start -nodes node_names
```



You can only execute this command at the advanced privilege level:
set -privilege advanced

You will see the following output:

Warning:

1. Current ONTAP version is 9.x

Before starting controller replacement operation, ensure that the new controllers are running the version 9.x

2. Verify that NVMEM or NVRAM batteries of the new nodes are charged, and charge them if they are not. You need to physically check the new nodes to see if the NVMEM or NVRAM batteries are charged. You can check the battery status either by connecting to a serial console or using SSH, logging into the Service Processor (SP) or Baseboard Management Controller (BMC) 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.

3. If a controller was previously part of a different cluster, run wipeconfig before using it as the replacement controller.

Do you want to continue? {y|n}: y

2. Press **y**, you will see the following output:

```
Controller replacement operation: Prechecks in progress.  
Controller replacement operation has been paused for user intervention.
```

The system runs the following prechecks; record the output of each precheck for use later in the procedure:

Precheck	Description
Cluster Health Check	Checks all the nodes in the cluster to confirm they are healthy.
MCC Cluster Check	Checks if the system is a MetroCluster configuration. The operation automatically detects if it is a MetroCluster configuration or not and performs the specific prechecks and verification checks. Only 4-node MetroCluster FC configuration is supported. In the case of 2-node MetroCluster configuration and 4-node MetroCluster IP configuration, the check fails. If the MetroCluster configuration is in switched over state, the check fails.
Aggregate Relocation Status Check	Checks whether an aggregate relocation is already in progress. If another aggregate relocation is in progress, the check fails.
Model Name Check	Checks whether the controller models are supported for this procedure. If the models are not supported, the task fails.
Cluster Quorum Check	Checks that the nodes being replaced are in quorum. If the nodes are not in quorum, the task fails.
Image Version Check	Checks that the nodes being replaced run the same version of ONTAP. If the ONTAP image versions are different, the task fails. 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. For instructions on how to upgrade ONTAP, refer to References to link to <i>Upgrade ONTAP</i> .
HA Status Check	Checks if both the nodes being replaced are in a high-availability (HA) pair configuration. If storage failover is not enabled for the controllers, the task fails.
Aggregate Status Check	If the nodes being replaced own aggregates for which they are not the home owner, the task fails. The nodes should not own any non-local aggregates.
Disk Status Check	If any nodes being replaced have missing or failed disks, the task fails. If any disks are missing, refer to References to link to <i>Disk and aggregate management with the CLI</i> , <i>Logical storage management with the CLI</i> , and <i>HA pair management</i> to configure storage for the HA pair.

Precheck	Description
Data LIF Status Check	Checks if any of the nodes being replaced have non- local data LIFs. The nodes should not contain any data LIFs for which they are not the home owner. If one of the nodes contains non-local data LIFs, the task fails.
Cluster LIF Status	Checks whether the cluster LIFs are up for both nodes. If the cluster LIFs are down, the task fails.
ASUP Status Check	If ASUP notifications are not configured, the task fails. You must enable ASUP before beginning the controller replacement procedure.
CPU Utilization Check	Checks if the CPU utilization is more than 50% for any of the nodes being replaced. If the CPU usage is more than 50% for a considerable period of time, the task fails.
Aggregate Reconstruction Check	Checks if reconstruction is occurring on any data aggregates. If aggregate reconstruction is in progress, the task fails.
Node Affinity Job Check	Checks if any node affinity jobs are running. If node affinity jobs are running, the check fails.

3. After the controller replacement operation is started and the prechecks are completed, the operation pauses enabling you to collect output information that you might need later when configuring node3.
4. Run the below set of commands as directed by the controller replacement procedure on the system console.

From the serial port connected to each node, run and save the output of the following commands individually:

```

° vservers services name-service dns show
° network interface show -curr-node local -role cluster,intercluster,node-
  mgmt,clustermgmt, data
° network port show -node local -type physical
° service-processor show -node local -instance
° network fcp adapter show -node local
° network port ifgrp show -node local
° network port vlan show
° system node show -instance -node local
° run -node local sysconfig
° storage aggregate show -node local
° volume show -node local
° network interface failover-groups show
° storage array config show -switch switch_name

```

- `system license show -owner local`

- `storage encryption disk show`



If NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE) using Onboard Key Manager is in use, keep the key manager passphrase ready to complete the key manager resync later in the procedure.

5. If your system uses self-encrypting drives, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use on the HA pair that you are upgrading. ONTAP software supports two types of self-encrypting drives:

- FIPS-certified NetApp Storage Encryption (NSE) SAS or NVMe drives
- Non-FIPS self-encrypting NVMe drives (SED)



You cannot mix FIPS drives with other types of drives on the same node or HA pair.

You can mix SEDs with non-encrypting drives on the same node or HA pair.

[Learn more about supported self-encrypting drives.](#)

Correct aggregate ownership if an ARL precheck fails

If the Aggregate Status Check fails, you must return aggregates owned by the partner node to the home owner node and initiate the precheck process again.

Steps

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

```
storage aggregate relocation start -node source_node -destination destination-  
node -aggregate-list *
```

2. 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
```

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.
```

After you finish

You must restart the controller replacement process:

```
system controller replace start -nodes node_names
```

License

Some features require 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.

Refer to [References](#) to link to the *NetApp Support Site* where you can obtain new 28-character license keys for ONTAP. 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, refer to [References](#) to link to the *System Administration Reference*.

Manage storage encryption using the Onboard Key Manager

You can use the Onboard Key Manager (OKM) to manage encryption keys. If you have the OKM set up, you must record the passphrase and backup material before beginning the upgrade.

Steps

1. Record the cluster-wide passphrase.

This is the passphrase that was entered when the OKM was configured or updated using the CLI or REST API.

2. Back up the key-manager information by running the `security key-manager onboard show-backup` command.

Quiesce the SnapMirror relationships (optional)

Before continuing with the procedure, you must confirm that all the SnapMirror relationships are quiesced. When a SnapMirror relationship is quiesced, it remains quiesced across reboots and failovers.

Steps

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

```
snapmirror show
```



If the status is "Transferring", you must abort those transfers:

```
snapmirror abort -destination-vserver vserver_name
```

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

2. Quiesce all relationships between the cluster:

```
snapmirror quiesce -destination-vserver *
```

Stage 2. Relocate and retire node1

Relocate non-root aggregates and NAS data LIFs owned by node1 to node2

Before you can replace node1 with node3, you must move the non-root aggregates and NAS data LIFs from node1 to node2 before eventually moving node1's resources to node3.

Before you begin

The operation must already be paused when you begin the task; you must manually resume the operation.

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.



The home owner for the aggregates and LIFs is not modified; only the current owner is modified.

Steps

1. Resume the aggregate relocation and NAS data LIF move operations:

```
system controller replace resume
```

All the non-root aggregates and NAS data LIFs are migrated from node1 to node2.

The operation pauses to enable you to verify whether all node1 non-root aggregates and non-SAN data LIFs have been migrated to node2.

2. Check the status of the aggregate relocation and NAS data LIF move operations:

```
system controller replace show-details
```

3. With the operation still paused, verify that all the non-root aggregates are online for their state on node2:

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

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,normal							
aggr_2	825.0GB	825.0GB	0%	online	1	node2	
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 <aggregate_name>
```

4. Verify that all the volumes are online on node2 by using 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 <vserver_name> -volume <volume_name>
```

The `vserver_name` to use with this command is found in the output of the previous `volume show` command.

5. If the ports currently hosting data LIFs will not exist on the new hardware, remove them from the broadcast domain:

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

6. 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-home-node  
nodename -status-admin up
```

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

- If you have not already saved them, record the VLAN and interface group information so you can re-create the VLANs and interface groups on node3 after node3 is booted up.

- b. Remove the VLANs from the interface groups:

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



Follow the corrective action to resolve any errors that are suggested by the vlan delete command.

- c. 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
```

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
```

- d. 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
```

Relocate failed or vetoed aggregates

If any aggregates fail to relocate or are vetoed, you must take manually relocate the aggregates, or override either the vetoes or destination checks, if necessary.

About this task

The system pauses the relocation operation due to the error.

Steps

1. Check the EMS logs to determine why the aggregate failed to relocate or was vetoed.
2. Relocate any failed or vetoed aggregates:

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

3. When prompted, enter *y*.
4. You can force relocation by using one of the following methods:

Option	Description
Overriding veto checks	Enter the following: storage aggregate relocation start -override -vetoes * -ndocontroller-upgrade true
Overriding destination checks	Enter the following: storage aggregate relocation start -overridedestination-checks * -ndo -controllerupgrade true

Retire node1

To retire node1, you resume the automated operation to disable the HA pair with node2 and shut node1 down correctly. Later in the procedure, you remove node1 from the rack or chassis.

Steps

1. Resume the operation:

```
system controller replace resume
```

2. Verify that node1 has been halted:

```
system controller replace show-details
```

After you finish

You can decommission node1 after the upgrade is completed. See [Decommission the old system](#).

Prepare for netboot

After you physically rack node3 and node4 later in the procedure, you might need to netboot them. The term "netboot" means you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you put a copy of the ONTAP 9 boot image onto a web server that the system can access.

Before you begin

- Verify that you can access a HTTP server with the system.
- Refer to [References](#) to link to the *NetApp Support Site* and download the necessary system files for your platform and the correct version of ONTAP.

About this task



You must netboot the new controllers if they do not have the same 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 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 [Stage 3 Installing and booting node3](#)

Steps

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

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

You will use the information in the directories in [Stage 3](#).

Stage 3. Install and boot node3

Install and boot node3

You must install node3 in the rack, transfer node1’s connections to node3, boot node3, and install ONTAP. You must then reassign any of node1’s spare disks, any disks belonging to the root volume, and any non-root aggregates that were not relocated to node2 earlier in the process, as outlined in this section.

About this task

The relocation operation is paused at the beginning of this stage. This process is largely automated; the operation pauses to enable you to check its status. You must manually resume the operation. In addition, you must verify the SAN LIFs have successfully moved to node3.

You need to netboot node3 if it does not have the same version of ONTAP 9 that is installed on node1. After you install node3, boot it 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, by following the instructions in [Prepare for netboot](#).



- For an AFF A800 or AFF C800 controller upgrade, you must ensure that all drives in the chassis are firmly seated against the midplane before removing node1. For more information, see [Replace the AFF A800 or AFF C800 controller modules](#).
- If you are upgrading a system with storage disks, you need to complete this entire section and then go to the [Configure FC ports on node3](#) and [Check and configure UTA/UTA2 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 the *Installation and Setup Instructions* for your node model.



If you're upgrading to a system with both nodes in the same chassis, install node4 and node3 in the chassis. If you don't install both nodes in the same chassis, when you boot node3, it behaves as if it were in a dual-chassis configuration, and when you boot node4, the interconnect between the nodes doesn't come up.

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

Cable the following connections using the *Installation and Setup Instructions* for the node3 platform, the appropriate disk shelf document, and the *HA pair management* documentation.

Refer to [References](#) to link to *HA pair management*.

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



You might 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 FC-VI cable connections from node1 to node3. If the new host does not have an FC-VI card, you might need to move the FC-VI card.

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 later.



When you boot node3, you might see the following warning 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.



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




Refer to [Prepare for netboot](#).

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



You must 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: <code>ifconfig e0M -auto</code>
Not running	<p>Manually configure the connection by entering the following command at the boot environment prompt:</p> <pre>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></pre> <p><i>filer_addr</i> is the IP address of the storage system (mandatory). <i>netmask</i> is the network mask of the storage system (mandatory). <i>gateway</i> is the gateway for the storage system. (mandatory). <i>dns_addr</i> is the IP address of a name server on your network (optional). <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> <div> Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</div>

7. Perform netboot on node3:

For...	Then...
FAS/AFF8000 series systems	<code>netboot http://<web_server_ip/path_to_web-accessible_directory>/netboot/kernel</code>
All other systems	<code>netboot http://<web_server_ip/path_to_web-accessible_directory>/<ontap_version>_image.tgz</code>

The `<path_to_the_web-accessible_directory>` should lead to where you downloaded the `<ontap_version>_image.tgz` in the section [Prepare for netboot](#).



Do not interrupt the boot.

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

This menu option downloads and installs the new 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 ONTAP, and not upgrades of controllers.



Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the incorrect image might install. This issue applies to all ONTAP releases. The netboot procedure combined with option (7) Install new software wipes the boot media and places the same ONTAP version on both image partitions.

9. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL:

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

10. Complete the following substeps to reboot the controller module:

- 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}
```

- b. Enter `y` to reboot 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 must be restored.

11. Select maintenance mode 5 from the boot menu and enter `y` when you are prompted to continue with the boot.
12. Verify that the controller and chassis are configured as ha:

```
ha-config show
```

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

```
Chassis HA configuration: ha
Controller HA configuration: ha
```



System records 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.

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

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

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

If you have a MetroCluster configuration, use the following commands to modify the controller and chassis:

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

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

14. Exit maintenance mode:

```
halt
```

Interrupt AUTOBOOT by pressing Ctrl-C at the boot environment prompt.

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

```
date
```

16. On node3, check the date by using the following command at the boot environment prompt:

```
show date
```

17. If necessary, set the date on node3:

```
set date mm/dd/yyyy
```

18. On node3, check the time by using the following command at the boot environment prompt:

```
show time
```

19. If necessary, set the time on node3:

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

20. In boot loader, set the partner system ID on node3:

```
setenv partner-sysid node2_sysid
```

For node3, partner-sysid must be that of node2.

- a. Save the settings:

```
saveenv
```

21. Verify the partner-sysid for node3:

```
printenv partner-sysid
```

22. If you have NetApp Storage Encryption (NSE) drives installed, perform the following steps.



If you have not already done so earlier in the procedure, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use.

- a. Set bootarg.storageencryption.support to true or false:

If the following drives are in use...	Then...
NSE drives that conform to FIPS 140-2 Level 2 self-encryption requirements	setenv bootarg.storageencryption.support true
NetApp non-FIPS SEDs	setenv bootarg.storageencryption.support false



You cannot mix FIPS drives with other types of drives on the same node or HA pair. You can mix SEDs with non-encrypting drives on the same node or HA pair.

- b. Contact NetApp Support for assistance with restoring the onboard key management information.

23. Boot the node into boot menu:

```
boot_ontap menu
```

What's next?

- If you have a system with an FC or UTA/UTA2 configuration, [set the FC or UTA/UTA2 configuration on node3](#).
- If you don't have an FC or UTA/UTA2 configuration, [reassign node1 disks to node3](#) so that node3 can recognize node1's disks.
- If you have a MetroCluster configuration, [reassign node1 disks to node3](#).

Set the FC or UTA/UTA2 configuration on node3

If node3 has onboard FC ports, onboard unified target adapter (UTA/UTA2) ports, or a UTA/UTA2 card, you must configure the settings before completing the rest of the procedure.

About this task

You might need to complete the section [Configure FC ports on node3](#), the section [Check and configure UTA/UTA2 ports on node3](#), or both sections.



NetApp marketing materials might use the term UTA2 to refer to converged network adapter (CNA) adapters and ports. However, the CLI uses the term CNA.

If node3 doesn't have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card, and you are upgrading a system with storage disks, you can skip to [Reassign node1 disks to node3](#).

Configure FC ports on node3

If node3 has FC ports, either onboard or on an add-on FC adapter, you must set port configurations on the node before you bring it into service because the ports are not preconfigured when the systems are shipped. If you don't configure the ports, 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 [Prepare the nodes for upgrade](#).

About this task

You can skip this section if your system does not have FC configurations. If your system has onboard UTA/UTA2 ports or a UTA/UTA2 card, you configure them in [Check and configure UTA/UTA2 ports on node3](#).



Enter the commands in this section at the Maintenance mode shell prompt.

Steps

1. Compare the FC settings on node3 with the settings that you captured earlier from node1.
2. Take one of the following actions to modify the FC ports on node3, as needed:

In Maintenance mode (option 5 at boot menu):

- To program as target ports:

```
ucadmin modify -m fc -t target <adapter>
```

For example: `ucadmin modify -m fc -t target 2a`

- To program initiator ports:

```
ucadmin modify -m fc -t initiator <adapter>
```

For example: `ucadmin modify -m fc -t initiator 2b`

3. Verify the new settings by using the following command and examining the output:

```
ucadmin show
```

4. Halt the node:

```
halt
```

5. Boot the system from LOADER prompt:

```
boot_ontap menu
```

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

7. Select option 5 from the boot menu for maintenance mode.

8. Perform one of the following actions:

- If node3 has a UTA/UTA2 card or UTA/UTA2 onboard ports, go to [Check and configure UTA/UTA2 ports on node3](#).
- If node3 doesn't have a UTA/UTA2 card or UTA/UTA2 onboard ports, skip [Check and configure UTA/UTA2 ports on node3](#) and go to [Reassign node1 disks to node3](#).

Check and configure UTA/UTA2 ports on node3

If node3 has onboard UTA/UTA2 ports or a UTA/UTA2 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 UTA/UTA2 ports.

About this task

If you want to use a Unified Target Adapter (UTA/UTA2) port for FC, you must first verify how the port is configured.



NetApp marketing materials might use the term UTA2 to refer to CNA adapters and ports. However, the CLI uses the term CNA.

You can use the `ucadmin show` command to view or verify the current port configuration, as shown in the following example output:

```
*> ucadmin show
```

Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
0e	fc	target	-	initiator	offline
0f	fc	target	-	initiator	offline
0g	fc	target	-	initiator	offline
0h	fc	target	-	initiator	offline
1a	fc	target	-	-	online
1b	fc	target	-	-	online

6 entries were displayed.

UTA/UTA2 ports can be configured into native FC mode or UTA/UTA2 mode. FC mode supports FC initiator and FC target; UTA/UTA2 mode allows concurrent NIC and FCoE traffic sharing the same 10GbE SFP+ interface and supports FC targets.

You might find UTA/UTA2 ports on an add-on adapter or on the controller motherboard, and have the following configurations, but you should check the configuration of the UTA/UTA2 ports on the node3 and change it, if

necessary:

- UTA/UTA2 cards ordered when the controller is ordered are configured before shipment to have the personality you request.
- UTA/UTA2 cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard UTA/UTA2 ports on new controllers are configured before shipment to have the personality you request.



You must be in Maintenance mode to configure UTA/UTA2 ports. Enter the commands in this section at the Maintenance mode shell prompt.

Steps

1. 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.

2. Verify the UTA/UTA2 port settings:

```
ucadmin show
```

Examine the output and determine whether the UTA/UTA2 ports have the personality you want.

The output in the following example shows that the type of adapter "1b" is changing to initiator and that the mode of adapters "2a" and "2b" is changing to "cna". The CNA mode allows you to use the card as a network adapter.

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

3. Take one of the following actions:

If the UTA/UTA2 ports...	Then...
Do not have the personality that you want	Go to Step 4 .
Have the personality that you want	Skip Step 4 through Step 8 and go to Step 9 .

4. Take one of the following actions:

If you are configuring...	Then...
Ports on a UTA/UTA2 card	Go to Step 5
Onboard UTA/UTA2 ports	Skip Step 5 and go to Step 6 .

- If the adapter is in initiator mode, and if the UTA/UTA2 port is online, take the UTA/UTA2 port offline:

```
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, change the configuration as needed:

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

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



You must use FC initiator for tape drives and MetroCluster configurations. You must use the FC target for SAN clients.

- Place any target ports online by entering the following command once for each port:

```
storage enable adapter <adapter_name>
```

- Cable the port.

- Exit maintenance mode:

```
halt
```

- Boot the node into boot menu by running `boot_ontap menu`.

What's next?

- If you are upgrading to an AFF A800 system, go to [Reassign node1 disks to node3, Step 9](#).
- For all other system upgrades, go to [Reassign node1 disks to node3, Step 1](#).

Reassign node1 disks to node3

You need to reassign the disks that belonged to node1 to node3 before verifying the node3 installation.

Steps

- Verify that node1 has stopped at the boot menu and reassign the disks of node1 to node3:

```
boot_after_controller_replacement
```

After a short delay, you are prompted to enter the name of the node that is being replaced. If there are shared disks (also called Advanced Disk Partitioning (ADP) or partitioned disks), you are prompted to enter the node name of the HA partner.

These prompts might get buried in the console messages. If you do not enter a node name or enter an incorrect name, you are prompted to enter the name again.

Expand the console output example

```
LOADER-A> boot_ontap menu

...
*****
*
* Press Ctrl-C for Boot Menu. *
*
*****

.
.
Please choose one of the following:
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
Selection (1-9)? 22/7

.
.
(boot_after_controller_replacement)  Boot after controller upgrade
(9a)                                  Unpartition all disks and
remove their ownership information.
(9b)                                  Clean configuration and
initialize node with partitioned disks.
(9c)                                  Clean configuration and
initialize node with whole disks.
(9d)                                  Reboot the node.
(9e)                                  Return to main boot menu.

Please choose one of the following:

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

```

.
This will replace all flash-based configuration with the last backup
to
disks. Are you sure you want to continue?: yes
.
.
Controller Replacement: Provide name of the node you would like to
replace: <name of the node being replaced>
Controller Replacement: Provide High Availability partner of node1:
<nodename of the partner of the node being replaced>
Changing sysid of node <node being replaced> disks.
Fetched sanown old_owner_sysid = 536953334 and calculated old sys id
= 536953334
Partner sysid = 4294967295, owner sysid = 536953334
.
.
.
Terminated
<node reboots>
.
.
System rebooting...
.
Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...
.
.
System rebooting...
.
.
.
WARNING: System ID mismatch. This usually occurs when replacing a
boot device or NVRAM cards!
Override system ID? {y|n} y
Login:
...

```

2. If the system goes into a reboot loop with the message `no disks found`, this is because it has reset the ports back to the target mode and therefore is unable to see any disks. Perform [Step 3](#) to [Step 8](#) on node3 to resolve this issue.
3. Press Ctrl-C during AUTOBOOT to stop the node at the `LOADER>` prompt.
4. At the `LOADER` prompt, enter maintenance mode:

```
boot_ontap maint
```

5. In maintenance mode, display all the previously set initiator ports that are now in target mode:

```
ucadmin show
```

Change the ports back to initiator mode:

```
ucadmin modify -m fc -t initiator -f adapter name
```

6. Verify that the ports have been changed to initiator mode:

```
ucadmin show
```

7. Exit maintenance mode:

```
halt
```



If you are upgrading from a system that supports external disks to a system that also supports external disks, go to [Step 8](#).

If you are upgrading from a system that supports external disks to a system that supports both internal and external disks, for example, an AFF A800 system, go to [Step 9](#).

8. At the LOADER prompt, boot up:

```
boot_ontap menu
```

Now, on booting, the node can detect all the disks that were previously assigned to it and can boot up as expected.

When the cluster nodes you are replacing use root volume encryption, ONTAP is unable to read the volume information from the disks. Restore the keys for the root volume:

- a. Return to the special boot menu:

```
LOADER> boot_ontap menu
```

Please choose one of the following:

- (1) Normal Boot.
- (2) Boot without /etc/rc.
- (3) Change password.
- (4) Clean configuration and initialize all disks.
- (5) Maintenance mode boot.
- (6) Update flash from backup config.
- (7) Install new software first.
- (8) Reboot node.
- (9) Configure Advanced Drive Partitioning.
- (10) Set Onboard Key Manager recovery secrets.
- (11) Configure node for external key management.

Selection (1-11)? 10

b. Select **(10) Set Onboard Key Manager recovery secrets**

c. Enter `y` at the following prompt:

```
This option must be used only in disaster recovery procedures. Are you sure?  
(y or n): y
```

d. At the prompt, enter the key-manager passphrase.

e. Enter the backup data when prompted.



You must have obtained the passphrase and backup data in the [Prepare the nodes for upgrade](#) section of this procedure.

f. After the system boots to the special boot menu again, run option **(1) Normal Boot**



You might encounter an error at this stage. If an error occurs, repeat the substeps in [Step 8](#) until the system boots normally.

9. If you are upgrading from a system with external disks to a system that supports internal and external disks (AFF A800 systems, for example), set the node1 aggregate as the root aggregate to confirm that node3 boots from the root aggregate of node1. To set the root aggregate, go to the boot menu on node3 and select option 5 to enter maintenance mode.



You must perform the following substeps in the exact order shown; failure to do so might cause an outage or even data loss.

The following procedure sets node3 to boot from the root aggregate of node1:

a. Enter maintenance mode:

```
boot_ontap maint
```

b. Check the RAID, plex, and checksum information for the node1 aggregate:

```
aggr status -r
```

c. Check the status of the node1 aggregate:

```
aggr status
```

d. If necessary, bring the node1 aggregate online:

```
aggr_online root_aggr_from_node1
```

e. Prevent the node3 from booting from its original root aggregate:

```
aggr offline root_aggr_on_node3
```

f. Set the node1 root aggregate as the new root aggregate for node3:

```
aggr options aggr_from_node1 root
```

- g. Verify that the root aggregate of node3 is offline and the root aggregate for the disks brought over from node1 is online and set to root:

```
aggr status
```



Failing to perform the previous substep might cause node3 to boot from the internal root aggregate, or it might cause the system to assume a new cluster configuration exists or prompt you to identify one.

The following shows an example of the command output:

```
-----  
Aggr              State    Status              Options  
  
aggr0_nst_fas8080_15 online  raid_dp, aggr      root, nosnap=on  
                                     fast zeroed  
                                     64-bit  
  
aggr0              offline raid_dp, aggr      diskroot  
                                     fast zeroed  
                                     64-bit  
-----
```

Map ports from node1 to node3

You must verify that the physical ports on node1 map correctly to the physical ports on node3, which will enable node3 to communicate with other nodes in the cluster and with the network after the upgrade.

About this task

Refer to [References](#) to link to the *Hardware Universe* to capture information about the ports on the new nodes. You will use the information later in this section.

Port settings might vary, depending on the model of the nodes. You must make the port and LIF configuration on the original node compatible with the planned use and configuration of the new node. 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 might not automatically rejoin quorum when it is rebooted until you change the software configuration 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 table:

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

2. Record all the cabling information for node3, the ports, broadcast domains, and IPspaces in the table.
3. Follow these steps to verify if the setup is a two-node switchless cluster:
 - a. Set the privilege level to advanced:

```
cluster::> set -privilege advanced
```

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

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

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

Enable Switchless Cluster: false/true
```

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

- c. Return to the administration privilege level:

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

cluster::>
```

4. Follow these steps to place node3 into quorum:

- a. Boot node3. See [Install and boot node3](#) 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
```

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	e0a	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
```

This example adds Cluster port "e1b" on node3:

```
network port modify -node node3 -port e1b -ip-space Cluster -mtu 9000
```

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

```
network port modify -node -port -ip-space Cluster -mtu 9000
```

This example adds Cluster port "e1b" on node4:

```
network port modify -node node4 -port e1b -ip-space Cluster -mtu 9000
```

- e. 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
```

- 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
```

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
```

5. Adjust the broadcast domains hosting your cluster LIFs and node-management/clustermanagement LIFs. Confirm 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 might need to migrate and modify the LIFs as follows:

- 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 vsServer -lif lif_name -home-port port_name
```

6. Adjust the broadcast domain membership of network ports used for intercluster LIFs using the same commands shown in [Step 5](#).
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 node1 that no longer exist on node3, follow these steps to delete them:

- a. Access the advanced privilege level on either node:

```
set -privilege advanced
```

- b. To delete the ports:

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

- c. 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
```

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-  
domainwide -failover-group fg1
```

Refer to [References](#) to link to *Network Management* or the *ONTAP 9 Commands: Manual Page Reference* for more information.

10. Verify the changes on node3:

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

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

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

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

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

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

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

Repeat Step 11 to verify that the cluster LIF is now listening on port 7700.

Join the quorum when a node has a different set of network ports

The node with the new controller boots and attempts to join the cluster automatically at first; however, if the new node has a different set of network ports, you must perform the following steps to confirm that the node successfully joins the quorum.

About this task

You can use these instructions for any relevant node. Node3 is used throughout the following sample.

Steps

1. Verify that the new cluster ports are in the Cluster broadcast domain by entering the following command

and checking its output:

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

The following example shows that port "e1a" is in the Cluster domain on node3:

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

2. Add the correct ports to the Cluster broadcast domain by entering the following command and checking its output:

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

This example adds Cluster port "e1b" on node3:

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

3. Migrate the cluster LIFs to the new ports, once for each LIF, using the following command:

```
network interface migrate -vserver Cluster -lif lif_name -source-node node3 -  
destination-node node3 -destination-port port_name
```

4. Modify the home port of the cluster LIFs:

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

5. If the cluster ports are not in the Cluster broadcast-domain, add them by using the following command:

```
network port broadcast-domain add-ports -ipSPACE Cluster -broadcast-domain  
Cluster - ports node:port
```

6. Remove the old ports from the Cluster broadcast domain. You can use for any relevant node. The following command removes port "e0d" on node3:

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

7. Verify the node has rejoined quorum:

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

8. Adjust the broadcast domains hosting your cluster LIFs and node-management/cluster management LIFs. Confirm 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 might need to migrate and modify the LIFs as follows:

- 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-port
```

- d. Modify a home port of a LIF:

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

Adjust the intercluster broadcast domains and migrate the intercluster LIFs, if necessary. The data LIFs remain unchanged.

Verify the node3 installation

After you install and boot node3, you must verify that it is installed correctly. You must wait for node3 to join quorum and then resume the relocation operation.

About this task

At this point in the procedure, the operation will have paused as node3 joins quorum.

Steps

1. Verify that node3 has joined quorum:

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

2. Verify that node3 is part of the same cluster as node2 and that it is healthy:

```
cluster show
```

3. Check the status of the operation and verify that the configuration information for node3 is the same as node1:

```
system controller replace show-details
```

If the configuration is different for node3, a system disruption might occur later in the procedure.

4. Check that the replaced controller is configured correctly for the MetroCluster configuration, the MetroCluster configuration should be in healthy state and not in switch over mode. Refer to [Verify the health of the MetroCluster configuration](#).

Re-create VLANs, interface groups, and broadcast domains on node3

After you confirm that node3 is in quorum and can communicate with node2, you must re-create node1's VLANs, interface groups, and broadcast domains on node3. You must also add the node3 ports to the newly re-created broadcast domains.

About this task

For more information on creating and re-creating VLANs, interface groups, and broadcast domains, go to [References](#) and link to *Network Management*.

Steps

1. Re-create the VLANs on node3 using the node1 information recorded in the [Relocate non-root aggregates and NAS data LIFs owned by node1 to node2](#) section:

```
network port vlan create -node node_name -vlan vlan-names
```

2. Re-create the interface groups on node3 using the node1 information recorded in the [Relocate non-root aggregates and NAS data LIFs owned by node1 to node2](#) section:

```
network port ifgrp create -node node_name -ifgrp port_ifgrp_names-distr-func
```

3. Re-create the broadcast domains on node3 using the node1 information recorded in the [Relocate non-root aggregates and NAS data LIFs owned by node1 to node2](#) section:

```
network port broadcast-domain create -ipSpace Default -broadcast-domain  
broadcast_domain_names -mtu mtu_size -ports  
node_name:port_name,node_name:port_name
```

4. Add the node3 ports to the newly re-created broadcast domains:

```
network port broadcast-domain add-ports -broadcast-domain  
broadcast_domain_names -ports node_name:port_name,node_name:port_name
```

Restore key-manager configuration on node3

If you are using NetApp Aggregate Encryption (NAE) or NetApp Volume Encryption (NVE) to encrypt volumes on the system you are upgrading, the encryption configuration must be synchronized to the new nodes. If you do not restore key-manager, when you relocate the node1 aggregates from node2 to node3 by using ARL, encrypted volumes will be taken offline.

Steps

1. To synchronize encryption configuration for Onboard Key Manager, run the following command at the cluster prompt:

For this ONTAP version...	Use this command...
ONTAP 9.6 or 9.7	<code>security key-manager onboard sync</code>
ONTAP 9.5	<code>security key-manager setup -node <i>node_name</i></code>

2. Enter the cluster-wide passphrase for the Onboard Key Manager.

Move non-root aggregates and NAS data LIFs owned by node1 from node2 to node3

After you verify the node3 installation and before you relocate aggregates from node2 to node3, you must move the NAS data LIFs belonging to node1 that are currently on node2 from node2 to node3. You also must verify that the SAN LIFs exist 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. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

- Cluster quorum check
- System ID check
- Image version check
- Target platform check
- Network reachability check

The operation pauses at this stage in the network reachability check.

2. Manually verify that the network and all VLANs, interface groups, and broadcast domains have been configured correctly.
3. Resume the relocation operation:

```
system controller replace resume
```

To complete the "Network Reachability" phase, ONTAP network configuration must be manually adjusted to match the new physical network configuration of the hardware. This includes assigning network ports to the correct broadcast domains, creating any required ifgrps and VLANs, and modifying the home-port parameter of network interfaces to the appropriate ports. Refer to the "Using aggregate relocation to upgrade controller hardware on a pair of nodes running ONTAP 9.x" documentation, Stages 3 and 5. Have all of these steps been manually completed? [y/n]

4. Enter `y` to continue.
5. The system performs the following checks:

- Cluster health check
- Cluster LIF status check

After performing these checks, the system relocates the non-root aggregates and NAS data LIFs owned by node1 to the new controller, node3.

The system pauses once the resource relocation is complete.

6. Check the status of the aggregate relocation and NAS data LIF move operations:

```
system controller replace show-details
```

7. Verify that the non-root aggregates and NAS data LIFs have been successfully relocated to node3.

If any aggregates fail to relocate or are vetoed, you must manually relocate the aggregates, or override either the vetoes or destination checks, if necessary. See [Relocate failed or vetoed aggregates](#) for more information.

8. Verify 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|fc -home-node node3
```

The system returns output similar to the following example:

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

	Logical	Status	Network	Current	Current	Is
Vserver	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 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:

```
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:

```
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:

```
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
```



You must confirm that you moved SAN LIFs to a port that has the same link speed as the original port.

- c. Modify the status of all LIFs to "up" so the LIFs can accept and send traffic on the node:

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

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

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

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

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

9. Resume the operation to prompt the system to perform the required post-checks:

```
system controller replace resume
```

The system performs the following post-checks:

- Cluster quorum check
- Cluster health check
- Aggregates reconstruction check
- Aggregate status check
- Disk status check
- Cluster LIF status check

Stage 4. Relocate and retire node2

Relocate non-root aggregates and NAS data LIFs from node2 to node3

Before replacing node2 with node4, you relocate the non-root aggregates and NAS data LIFs that are owned by node2 to node3.

Before you begin

After the post-checks from the previous stage complete, the resource release for node2 starts automatically. The non-root aggregates and non-SAN data LIFs are migrated from node2 to 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.

After the aggregates and LIFs are migrated, the operation is paused for verification purposes. At this stage, you must verify whether or not all the non-root aggregates and non-SAN data LIFs are migrated to node3.



The home owner for the aggregates and LIFs are not modified; only the current owner is modified.

Steps

1. Verify that all the non-root aggregates are online and their state on node3:

```
storage aggregate show -node <node3> -state online -root false
```

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

```
cluster::> storage aggregate show -node node3 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	normal					
aggr_2	825.0GB	825.0GB	0%	online	1	node2
raid_dp	normal					

2 entries were displayed.

If the aggregates have gone offline or become foreign on node3, bring them online by using the following command on node3, once for each aggregate:

```
storage aggregate online -aggregate <aggregate_name>
```

2. Verify that all the volumes are online on node3 by using the following command on node3 and examining the output:

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

If any volumes are offline on node3, bring them online by using the following command on node3, once for each volume:

```
volume online -vserver <vserver_name> -volume <volume_name>
```

The `vserver_name` to use with this command is found in the output of the previous `volume show` command.

3. Verify that the LIFs have been moved to the correct ports and have a status of `up`. 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> -home-node
<node_name> -status-admin up
```

4. If the ports currently hosting data LIFs will not exist on the new hardware, remove them from the broadcast domain:

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

5. Verify that there are no data LIFs remaining on node2 by entering the following command and examining the output:

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

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

- a. Record VLAN and interface group information so you can re-create the VLANs and interface groups on node3 after node3 is booted up.
- b. Remove the VLANs from the interface groups:

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

- c. Check if there are any interface groups configured on the node by entering the following command and examining its output:

```
network port ifgrp show -node node2 -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: node3
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
```

- d. 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
```

Retire node2

To retire node2, you shut down node2 correctly and then remove it from the rack or chassis.

Steps

1. Resume the operation:

```
system controller replace resume
```

The node halts automatically.

After you finish

You can decommission node2 after the upgrade is completed. See [Decommission the old system](#).

Stage 5. Install and boot node4

Install and boot node4

You must install node4 in the rack, transfer the node2 connections to to node4, boot node4, and install ONTAP. You must then reassign any spare disks on node2, any disks belonging to the root volume, and any non-root aggregates that were not relocated to node3 earlier in the process, as outlined in this section.

About this task

The relocation operation is paused at the beginning of this stage. This process is mostly automated; the operation pauses to enable you to check its status. You must manually resume the operation. In addition, you must verify that the NAS data LIFs have successfully moved to node4.

You need to netboot node4 if the ONTAP version on node4 is different to the ONTAP version on node2. After you install node4, boot it 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 by following the instructions in [Prepare for netboot](#).



- For an AFF A800 or AFF C800 controller upgrade, you must ensure that all drives in the chassis are firmly seated against the midplane before removing node2. For more information, see [Replace the AFF A800 or AFF C800 controller modules](#).
- If you are upgrading a system with storage disks, you must complete this entire section and then proceed to [Set the FC or UTA/UTA2 configuration on node4](#), entering commands at the cluster prompt.

Steps

1. Make sure that node4 has sufficient rack space.

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

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

Cable the following connections, using the *Installation and Setup Instructions* for the node4 platform, the appropriate disk shelf document, and the *HA pair management* documentation.

Refer to [References](#) to link to *HA pair management*.

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



You might 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.
For the MetroCluster configuration, you must move the FC-VI cable connections from node2 to node4. If the new host does not have an FC-VI card, you might need to move the FC-VI card.

4. 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.



When you boot node4, you might see the following warning 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.



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



Refer to [Prepare for netboot](#).

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



You must 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: <code>ifconfig e0M -auto</code>
Not running	<p>Manually configure the connection by entering the following command at the boot environment prompt:</p> <pre>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></pre> <p><i>filer_addr</i> is the IP address of the storage system (mandatory). <i>netmask</i> is the network mask of the storage system (mandatory). <i>gateway</i> is the gateway for the storage system (mandatory). <i>dns_addr</i> is the IP address of a name server on your network (optional). <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>NOTE: Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p>

7. Perform netboot on node4:

For...	Then...
FAS/AFF8000 series systems	<code>netboot http://<web_server_ip/path_to_web-accessible_directory>/netboot/kernel</code>
All other systems	<code>netboot http://<web_server_ip/path_to_web-accessible_directory>/<ontap_version>_image.tgz</code>

The `<path_to_the_web-accessible_directory>` should lead to where you downloaded the `<ontap_version>_image.tgz` in Step 1 in the section [Prepare for netboot](#).



Do not interrupt the boot.

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

This menu option downloads and installs the new 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 ONTAP, and not upgrades of controllers.



Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the incorrect image might install. This issue applies to all ONTAP releases. The netboot procedure combined with option (7) `Install new software` wipes the boot media and places the same ONTAP version on both image partitions.

9. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL:

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

10. Complete the following substeps to reboot the controller module:

- 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}
```

- 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 must be restored.

11. Select maintenance mode 5 from the boot menu and enter `y` when you are prompted to continue with the boot.
12. Verify that the controller and chassis are configured as HA:

```
ha-config show
```

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

```
Chassis HA configuration: ha  
Controller HA configuration: ha
```



System records 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.

13. If the controller and chassis are not configured as HA, use the following commands to correct the configuration:

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

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

If you have a MetroCluster configuration, use the following commands to modify the controller and chassis:

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

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

14. Exit maintenance mode:

```
halt
```

Interrupt AUTOBOOT by pressing Ctrl-C at the boot environment prompt.

15. On node3, check the system date, time, and time zone:

```
date
```

16. On node4, check the date by using the following command at the boot environment prompt:

```
show date
```

17. If necessary, set the date on node4:

```
set date mm/dd/yyyy
```

18. On node4, check the time by using the following command at the boot environment prompt:

```
show time
```

19. If necessary, set the time on node4:

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

20. In boot loader, set the partner system ID on node4:

```
setenv partner-sysid node3_sysid
```

For node4, partner-sysid must be that of node3.

Save the settings:

```
saveenv
```

21. Verify the partner-sysid for node4:

```
printenv partner-sysid
```

22. If you have NetApp Storage Encryption (NSE) drives installed, perform the following steps:



If you have not already done so earlier in the procedure, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use.

a. Set `bootarg.storageencryption.support` to true or false:

If the following drives are in use...	Then...
NSE drives that conform to FIPS 140-2 Level 2 self-encryption requirements	<code>setenv bootarg.storageencryption.support true</code>
NetApp non-FIPS SEDs	<code>setenv bootarg.storageencryption.support false</code>



You cannot mix FIPS drives with other types of drives on the same node or HA pair. You can mix SEDs with non-encrypting drives on the same node or HA pair.

b. Contact NetApp Support for assistance with restoring the onboard key management information.

23. Boot the node into boot menu:

```
boot_ontap menu
```

What's next?

- If you have a system with an FC or UTA/UTA2 configuration, [set the FC or UTA/UTA2 configuration on node4](#).
- If you don't have an FC or UTA/UTA2 configuration, [reassign node2 disks to node4, Step 1](#) so that node4 can recognize node2's disks.
- If you have a MetroCluster configuration, [set the FC or UTA/UTA2 configuration on node4](#) to detect the disks attached to the node.

Set the FC or UTA/UTA2 configuration on node4

If node4 has onboard FC ports, onboard unified target adapter (UTA/UTA2) ports, or a UTA/UTA2 card, you must configure the settings before completing the rest of the procedure.

About this task

You might need to complete [Configure FC ports on node4](#) or [Check and configure UTA/UTA2 ports on node4](#), or both sections.



If node4 doesn't have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card, and you are upgrading a system with storage disks, you can skip to [Reassign node2 disks to node4](#).

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

Configure FC ports on node4

If node4 has FC ports, either onboard or on an add-on FC adapter, you must set port configurations on the node before you bring it into service because the ports are not preconfigured when the systems are shipped. If you don't configure the ports as required, 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 [Prepare the nodes](#)

for upgrade.

About this task

You can skip this section if your system does not have FC configurations. If your system has onboard UTA/UTA2 ports or a UTA/UTA2 adapter, you configure them in [Check and configure UTA/UTA2 ports on node4](#).



Enter the commands in this section at the Maintenance mode shell prompt.

Steps

1. Display information about all FC and converged network adapters on the system:

```
system node hardware unified-connect show
```

2. Compare the FC settings on node4 with the settings that you captured earlier from node1.
3. Modify the FC ports on node4 as needed:

- To program as target ports:

```
ucadmin modify -m fc -t target adapter
```

For example: `ucadmin modify -m fc -t target 2a`

- To program initiator ports:

```
ucadmin modify -m fc -t initiator adapter
```

`-t` is the FC4 type: target or initiator.

For example: `ucadmin modify -m fc -t initiator 2b`

4. Halt the node:

```
halt
```

5. Boot the system from LOADER prompt:

```
boot_ontap menu
```

6. After you enter the command, wait until the system stops at the boot environment prompt.
7. Select option 5 from the boot menu for maintenance mode.
8. Take one of the following actions:
 - Go to [Check and configure UTA/UTA2 ports on node4](#) if node4 has a UTA/UTA2 card or UTA/UTA2 onboard ports.
 - If node4 doesn't have a UTA/UTA2 card or UTA/UTA2 onboard ports, skip *Check and configure UTA/UTA2 ports on node4* and go to [Reassign node2 disks to node4](#).

Check and configure UTA/UTA2 ports on node4

If node4 has onboard UTA/UTA2 ports or a UTA/UTA2A 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 UTA/UTA2 ports.

About this task

UTA/UTA2 ports can be configured into native FC mode or UTA/UTA2A mode. FC mode supports FC initiator and FC target; UTA/UTA2 mode allows concurrent NIC and FCoE traffic to share the same 10GbE SFP+ interface and supports FC target.



NetApp marketing materials might use the term UTA2 to refer to CNA adapters and ports. However, the CLI uses the term CNA.

UTA/UTA2 ports might be on an adapter or on the controller with the following configurations:

- UTA/UTA2 cards ordered at the same time as the controller are configured before shipment to have the personality you requested.
- UTA/UTA2 cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard UTA/UTA2 ports on new controllers are configured (before shipment) to have the personality you requested.

However, you should check the configuration of the UTA/UTA2 ports on node4 and change it, if necessary.



Enter the commands in this section at the Maintenance mode shell prompt.

Steps

1. Check how the ports are currently configured on node4:

```
system node hardware unified-connect show
```

The system displays output similar to the following example:

```
*> ucadmin show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
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

```
*>
```

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. Verify the settings:

```
ucadmin show
```

Examine the output of the `ucadmin show` command and determine whether the UTA/UTA2 ports have the personality you want.

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`:

```
*> 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.
*>
```

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 9 and go to Step 10 .

5. Take one of the following actions:

If you are configuring...	Then...
Ports on a UTA/UTA2 card	Go to Step 6
Onboard UTA/UTA2 ports	Skip Step 6 and go to Step 7 .

6. If the adapter is in initiator mode, and if the UTA/UTA2 port is online, take the UTA/UTA2 port offline:

```
storage disable adapter adapter_name
```

Adapters in target mode are automatically offline in Maintenance mode.

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

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

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



You must use FC initiator for tape drives and MetroCluster configurations. You must use the FC target for SAN clients.

8. Place any target ports online by entering the following command, once for each port:

```
storage enable adapter <adapter_name>
```

9. Cable the port.

10. Exit Maintenance mode:

```
halt
```

11. Boot the node into boot menu:

```
boot_ontap menu
```

What's next?

- If you are upgrading to an AFF A800 system, go to [Reassign node2 disks to node4, Step 9](#).
- For all other system upgrades, go to [Reassign node2 disks to node4, Step 1](#).

Reassign node2 disks to node4

You need to reassign the disks that belonged to node2 to node4 before verifying the node4 installation..

Steps

1. Verify that node2 has stopped at the boot menu and reassign the disks of node2 to node4:

```
boot_after_controller_replacement
```

After a short delay, you are prompted to enter the name of the node that is being replaced. If there are shared disks (also called Advanced Disk Partitioning (ADP) or partitioned disks), you are prompted to enter the node name of the HA partner.

These prompts might get buried in the console messages. If you do not enter a node name or enter an incorrect name, you are prompted to enter the name again.

Expand the console output example

```
LOADER-A> boot_ontap menu ...
*****
*
* Press Ctrl-C for Boot Menu. *
*
*****

.
.
Please choose one of the following:

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
Selection (1-9)? 22/7
.
.
(boot_after_controller_replacement) Boot after controller upgrade
(9a) Unpartition all disks and remove
their ownership information.
(9b) Clean configuration and
initialize node with partitioned disks.
(9c) Clean configuration and
initialize node with whole disks.
(9d) Reboot the node.
(9e) Return to main boot menu.

Please choose one of the following:

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

```

.
This will replace all flash-based configuration with the last backup
to disks. Are you sure you want to continue?: yes
.
.
Controller Replacement: Provide name of the node you would like to
replace: <name of the node being replaced>
Controller Replacement: Provide High Availability partner of node1:
<nodename of the partner of the node being replaced>
Changing sysid of node <node being replaced> disks.
Fetched sanown old_owner_sysid = 536953334 and calculated old sys id
= 536953334
Partner sysid = 4294967295, owner sysid = 536953334
.
.
.
Terminated
<node reboots>
.
.
System rebooting...
.
Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...
.
.
System rebooting...
.
.
.
WARNING: System ID mismatch. This usually occurs when replacing a
boot device or NVRAM cards!
Override system ID? {y|n} y
Login: ...

```

2. If the system goes into a reboot loop with the message `no disks found`, this is because it has reset the ports back to the target mode and therefore is unable to see any disks. Perform [Step 3](#) through [Step 8](#) on node4 to resolve this issue.
3. Press Ctrl-C during AUTOBOOT to stop the node at the `LOADER>` prompt.
4. At the `LOADER` prompt, enter maintenance mode:

```
boot_ontap maint
```

5. In maintenance mode, display all the previously set initiator ports that are now in target mode:

```
ucadmin show
```

Change the ports back to initiator mode:

```
ucadmin modify -m fc -t initiator -f adapter name
```

6. Verify that the ports have been changed to initiator mode:

```
ucadmin show
```

7. Exit maintenance mode:

```
halt
```



If you are upgrading from a system that supports external disks to a system that also supports external disks, go to [Step 8](#).

If you are upgrading from a system that uses external disks to a system that supports both internal and external disks, for example, an AFF A800 system, go to [Step 9](#).

8. At the LOADER prompt, boot up:

```
boot_ontap menu
```

Now, on booting, the node can detect all the disks that were previously assigned to it and can boot up as expected.

When the cluster nodes you are replacing use root volume encryption, ONTAP is unable to read the volume information from the disks. Restore the keys for the root volume:

a. Return to the special boot menu:

```
LOADER> boot_ontap menu
```

Please choose one of the following:

- (1) Normal Boot.
- (2) Boot without /etc/rc.
- (3) Change password.
- (4) Clean configuration and initialize all disks.
- (5) Maintenance mode boot.
- (6) Update flash from backup config.
- (7) Install new software first.
- (8) Reboot node.
- (9) Configure Advanced Drive Partitioning.
- (10) Set Onboard Key Manager recovery secrets.
- (11) Configure node for external key management.

Selection (1-11)? 10

b. Select **(10) Set Onboard Key Manager recovery secrets**

c. Enter `y` at the following prompt:

```
This option must be used only in disaster recovery procedures. Are you sure?  
(y or n): y
```

d. At the prompt, enter the key-manager passphrase.

e. Enter the backup data when prompted.



You must have obtained the passphrase and backup data in the [Prepare the nodes for upgrade](#) section of this procedure.

f. After the system boots to the special boot menu again, run option **(1) Normal Boot**



You might encounter an error at this stage. If an error occurs, repeat the substeps in [Step 8](#) until the system boots normally.

9. If you are upgrading from a system with external disks to a system that supports internal and external disks (AFF A800 systems, for example), set the node2 aggregate as the root aggregate to confirm that node4 boots from the root aggregate of node2. To set the root aggregate, go to the boot menu on node4 and select option 5 to enter maintenance mode.



You must perform the following substeps in the exact order shown; failure to do so might cause an outage or even data loss.

The following procedure sets node4 to boot from the root aggregate of node2:

a. Enter maintenance mode:

```
boot_ontap maint
```

b. Check the RAID, plex, and checksum information for the node2 aggregate:

```
aggr status -r
```

c. Check the status of the node2 aggregate:

```
aggr status
```

d. If necessary, bring the node2 aggregate online:

```
aggr_online root_aggr_from_node2
```

e. Prevent the node4 from booting from its original root aggregate:

```
aggr offline root_aggr_on_node4
```

f. Set the node2 root aggregate as the new root aggregate for node4:

```
aggr options aggr_from_node2 root
```

Map ports from node2 to node4

You must verify that the physical ports on node2 map correctly to the physical ports on node4, which will enable node4 to communicate with other nodes in the cluster and with the network after the upgrade.

About this task

Refer to [References](#) to link to the *Hardware Universe* to capture information about the ports on the new nodes. You will use the information later in this section.

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 might not automatically rejoin the 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 node2 cabling information for node2, the ports, broadcast domains, and IPspaces, in the table:

LIF	Node2 ports	Node2 IPspaces	Node2 broadcast domains	Node4 ports	Node4 IPspaces	Node4 broadcast domains
Cluster 1						
Cluster 2						
Cluster 3						
Cluster 4						
Node management						
Cluster management						
Data 1						
Data 2						
Data 3						
Data 4						
SAN						
Intercluster port						

2. Record all the cabling information for node4, the ports, broadcast domains, and IPspaces, in the table.

3. Follow these steps to verify if the setup is a two-node switchless cluster:

a. Set the privilege level to advanced:

```
cluster::> set -privilege advanced
```

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

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

```
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. Follow these steps to place node4 into quorum:

a. Boot node4. See [Install and boot 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
```

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     e0a  Cluster
```

c. If the cluster ports are not in the Cluster broadcast-domain, add them with the following command:

```
broadcast-domain add-ports -ipspace Cluster -broadcast-domain Cluster -ports
node:port
```

d. Add the correct ports to the Cluster broadcast domain:

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

This example adds Cluster port "e1b" on node4:

```
network port modify -node node4 -port e1b -ipspace Cluster -mtu 9000
```

- 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
```

This command removes port "e0d" on node4:

```
network port broadcast-domain remove-ports -ipspace 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/clustermanagement LIFs. Confirm 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 vserver -lif lif_name -home-port port_name
```

6. Adjust the intercluster broadcast domains and migrate the intercluster LIFs, if necessary, using the same commands shown in [Step 5](#).
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
```

- b. To delete the ports:

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

- c. 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
```

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-domainwide -failover-group fg1
```

Refer to [References](#) to link to *Network Management* or the *ONTAP 9 Commands: Manual Page Reference* and see *Configuring failover settings on a LIF* for more information.

10. Verify the changes on node4:

```
network port show -node node4
```

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

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

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

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

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

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

Repeat step 11 to verify that the cluster LIF is now listening on port 7700.

Join the quorum when a node has a different set of network ports

The node with the new controller boots and attempts to join the cluster automatically at first; however, if the new node has a different set of network ports, you must perform the following steps to confirm that the node successfully joins the quorum.

About this task

You can use these instructions for any relevant node. Node3 is used throughout the following sample.

Steps

1. Verify that the new cluster ports are in the Cluster broadcast domain by entering the following command and checking the output:

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

The following example shows that port "e1a" is in the Cluster domain on node3:

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

2. Add the correct ports to the Cluster broadcast domain by entering the following command and checking the output:

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

This example adds Cluster port "e1b" on node3:

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

3. Migrate the cluster LIFs to the new ports, once for each LIF, using the following command:

```
network interface migrate -vserver Cluster -lif lif_name -source-node node3  
destination-node node3 -destination-port port_name
```

4. Modify the home port of the cluster LIFs as follows:

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

5. 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 -broadcastdomain  
Cluster ports node:port
```

6. Remove the old ports from the Cluster broadcast domain. You can use for any relevant node. The following command removes port "e0d" on node3:

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

7. Verify the node has rejoined quorum as follows:

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

8. Adjust the broadcast domains hosting your cluster LIFs and node-management/cluster management LIFs. Confirm 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 might need to migrate and modify the LIFs as follows:

- 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-port
```

- d. Modify a home port of a LIF:

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

Adjust the intercluster broadcast domains and migrate the intercluster LIFs, if necessary. The data LIFs remain unchanged.

Verify the node4 installation

After you install and boot node4, you must verify that it is installed correctly, that it is part of the cluster, and that it can communicate with node3.

About this task

At this point in the procedure, the operation will have paused as node4 joins quorum.

Steps

1. Verify that node4 has joined quorum:

```
cluster show -node node4 -fields health
```

2. Verify that node4 is part of the same cluster as node3 and healthy by entering the following command:

```
cluster show
```

3. Check the status of the operation and verify that the configuration information for node4 is the same as node2:

```
system controller replace show-details
```

If the configuration is different for node4, a system disruption might occur later in the procedure.

4. Check that the replaced controller is configured correctly for MetroCluster configuration and not in switch-over mode.



At this stage MetroCluster configuration will not be in a normal state and you might have errors to resolve. See [Verify the health of the MetroCluster configuration](#).

Re-create VLANs, interface groups, and broadcast domains on node4

After you confirm that node4 is in quorum and can communicate with node3, you must re-create node2's VLANs, interface groups, and broadcast domains on node4. You must also add the node3 ports to the newly re-created broadcast domains.

About this task

For more information on creating and re-creating VLANs, interface groups, and broadcast domains, go to [References](#) and link to *Network Management*.

Steps

1. Re-create the VLANs on node4 using the node2 information recorded in the [Relocate non-root aggregates and NAS data LIFs from node2 to node3](#) section:

```
network port vlan create -node node4 -vlan vlan-names
```

2. Re-create the interface groups on node4 using the node2 information recorded in the [Relocate non-root aggregates and NAS data LIFs from node2 to node3](#) section:

```
network port ifgrp create -node node4 -ifgrp port_ifgrp_names-distr-func
```

3. Re-create the broadcast domains on node4 using the node2 information recorded in the [Relocate non-root aggregates and NAS data LIFs from node2 to node3](#) section:

```
network port broadcast-domain create -ipspace Default -broadcast-domain  
broadcast_domain_names -mtu mtu_size -ports  
node_name:port_name,node_name:port_name
```

4. Add the node4 ports to the newly re-created broadcast domains:

```
network port broadcast-domain add-ports -broadcast-domain  
broadcast_domain_names -ports node_name:port_name,node_name:port_name
```

Restore key-manager configuration on node4

If you are using NetApp Aggregate Encryption (NAE) or NetApp Volume Encryption (NVE) to encrypt volumes on the system you are upgrading, the encryption configuration must be synchronized to the new nodes. If you do not restore key-manager, when you relocate the node2 aggregates from node3 to node4 by using ARL, encrypted volumes will be taken offline.

Steps

1. To synchronize encryption configuration for Onboard Key Manager, run the following command at the cluster prompt:

For this ONTAP version...	Use this command...
ONTAP 9.6 or 9.7	<code>security key-manager onboard sync</code>
ONTAP 9.5	<code>security key-manager setup -node <i>node_name</i></code>

2. Enter the cluster-wide passphrase for the Onboard Key Manager.

Move non-root aggregates and NAS data LIFs owned by node2 from node3 to node4

After you verify the node4 installation and before you relocate aggregates from node3 to node4, you must move the NAS data LIFs belonging to node2 that are currently on node3 from node3 to node4. You also need to verify the SAN LIFs exist 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 will verify that the LIFs are healthy and located on appropriate ports after you bring node4 online.

Steps

1. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

- Cluster quorum check
- System ID check
- Image version check
- Target platform check
- Network reachability check

The operation pauses at this stage in the network reachability check.

2. Manually verify that the network and all VLANs, interface groups, and broadcast domains have been configured correctly.
3. Resume the relocation operation:

```
system controller replace resume
```

To complete the "Network Reachability" phase, ONTAP network configuration must be manually adjusted to match the new physical network configuration of the hardware. This includes assigning network ports to the correct broadcast domains, creating any required ifgrps and VLANs, and modifying the home-port parameter of network interfaces to the appropriate ports. Refer to the "Using aggregate relocation to upgrade controller hardware on a pair of nodes running ONTAP 9.x" documentation, Stages 3 and 5. Have all of these steps been manually completed? [y/n]

4. Enter `y` to continue.

5. The system performs the following checks:

- Cluster health check
- Cluster LIF status check

After performing these checks, the system relocates the non-root aggregates and NAS data LIFs owned by node2 to the new controller, node4.

The system pauses once the resource relocation is complete.

6. Check the status of the aggregate relocation and NAS data LIF move operations:

```
system controller replace show-details
```

7. Manually verify that the non-root aggregates and NAS data LIFs have been successfully relocated to node4.

If any aggregates fail to relocate or are vetoed, you must manually relocate the aggregates, or override either the vetoes or destination checks, if necessary. See the section [Relocate failed or vetoed aggregates](#) for more information.

8. Confirm that the SAN LIFs are on the correct ports on node4 by completing the following substeps:

a. Enter the following command and examine its output:

```
network interface show -data-protocol iscsi|fc -home-node node4
```

The system returns output similar to the following example:


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

	Logical	Status	Network	Current	Current	Is
Vserver	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 or that need to be mapped to a different port, move them to an appropriate port on node4 by completing the following substeps:

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

```
network interface modify -vserver vs0 -lif lif1 -status
-admin down
```

- ii. Remove the LIF from the port set:

```
portset remove -vserver vs0 -portset portset1 -port-name
lif1
```

- iii. Enter one of the following commands:

- Move a single LIF by entering the following command:

```
network interface modify -vserver vs0 -lif lif1 -home
-port port2
```

- 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 port1 -home-node node1
-role data} -home-port port2 -home-node node2
```

- Add the LIFs back to the port set:

```
portset add -vserver vs0 -portset portset1 -port-name
lif1
```



You must confirm that you move SAN LIFs to a port that has the same link speed as the original port.

- c. 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 node4 -lif data
-statusadmin up
```

- d. 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 <node4> -role data
```

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

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

9. Resume the operation to prompt the system to perform the required post-checks:

```
system controller replace resume
```

The system performs the following post-checks:

- Cluster quorum check
- Cluster health check
- Aggregates reconstruction check
- Aggregate status check
- Disk status check
- Cluster LIF status check

Stage 6. Complete the upgrade

Manage authentication using KMIP servers

With ONTAP 9.5 to 9.7, you can use Key Management Interoperability Protocol (KMIP) servers to manage authentication keys.

Steps

1. Add a new controller:

```
security key-manager setup -node new_controller_name
```

2. Add the key manager:

```
security key-manager -add key_management_server_ip_address
```

3. Verify that the key management servers are configured and available to all nodes in the cluster:

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

4. Restore the authentication keys from all linked key management servers to the new node:

```
security key-manager restore -node new_controller_name
```

Confirm that the new controllers are set up correctly

To confirm 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 confirm that node3 owns node1's aggregates and that node4 owns node2's aggregates, and that the volumes for both nodes are online.

Steps

1. After the post-checks of node2, the storage failover and cluster HA pair for the node2 cluster are enabled. When the operation is done, both nodes show as completed and the system performs some cleanup operations.
2. Verify that storage failover is enabled:

```
storage failover show
```

The following example shows the output of the command when storage failover is enabled:

```
cluster::> storage failover show
```

Takeover			
Node	Partner	Possible	State Description
node3	node4	true	Connected to node4
node4	node3	true	Connected to node3

3. Verify that node3 and node4 belong to the same cluster by using the following command and examining the output:

```
cluster show
```

4. Verify that node3 and node4 can access each other's storage by using the following command and examining the output:

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

5. Verify that neither node3 nor node4 owns data LIFs home-owned by other nodes in the cluster by using the following command and examining the output:

```
network interface show
```

If neither node3 or node4 owns data LIFs home-owned by other nodes in the cluster, revert the data LIFs to their home owner:

```
network interface revert
```

6. Verify that node3 owns the aggregates from node1 and that node4 owns the aggregates from node2:

```
storage aggregate show -owner-name <node3>
```

```
storage aggregate show -owner-name <node4>
```

7. Determine whether any volumes are offline:

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

```
volume show -node <node4> -state offline
```

8. If any volumes are offline, compare them with the list of offline volumes that you captured in the section [Prepare the nodes for upgrade](#), and bring online any of the offline volumes, as required, by using the following command, once for each volume:

```
volume online -vserver <vserver_name> -volume <volume_name>
```

9. Install new licenses for the new nodes by using the following command 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 multiple licenses at once, separating each license key by a comma.

10. Remove all of the old licenses from the original nodes by using one of the following commands:

```
system license clean-up -unused -expired
```

```
system license delete -serial-number <node_serial_number> -package  
<licensable_package>
```

- Delete all expired licenses:

```
system license clean-up -expired
```

- Delete all unused licenses:

```
system license clean-up -unused
```

- Delete a specific license from a cluster by using 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.

11. Verify that the licenses are correctly installed by using the following command and examining the output:

```
system license show
```

You can compare the output with the output that you captured in the section [Prepare the nodes for upgrade](#).

12. If self-encrypting drives are being used in the configuration and you have set the `kmip.init.maxwait` variable to `off` (for example, in [Install and boot node4, Step 22](#)), you must unset the variable:

```
set diag; systemshell -node node_name -command sudo kenv -u -p  
kmip.init.maxwait
```

13. Configure the SPs by using the following command on both nodes:

```
system service-processor network modify -node node_name
```

Refer to [References](#) to link to the *System Administration Reference* for information about the SPs and the *ONTAP 9 Commands: Manual Page Reference* for detailed information about the `system service-processor network modify` command.

14. If you want to set up a switchless cluster on the new nodes, refer to [References](#) to link to the *NetApp Support Site* and follow the instructions in *Transitioning to a two-node switchless cluster*.

After you finish

If Storage Encryption is enabled on node3 and node4, complete the section [Set up Storage Encryption on the new controller module](#). Otherwise, complete the section [Decommission the old system](#).

Set 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. You must 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 the previous step 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
```

Set up NetApp Volume or Aggregate Encryption on the new controller module

If the replaced controller or high availability (HA) partner of the new controller uses NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE), you must configure the new controller module for NVE or NAE.

About this task

This procedure includes steps that are performed on the new controller module. You must enter the command on the correct node.

ONTAP 9.6 and 9.7

Configure NVE or NAE on controllers running ONTAP 9.6 or 9.7

Steps

1. Verify that the key management servers are still available, their status, and their authentication key information:

```
security key-manager key query -node node
```

2. Add the key management servers listed in the previous step 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.

- Restore authentication for external key manager:

```
security key-manager external restore
```

This command needs the Onboard Key Manager (OKM) passphrase.

For more information, see the Knowledge Base article [How to restore external key manager server configuration from the ONTAP boot menu](#).

- Restore authentication for the OKM:

```
security key-manager onboard sync
```

ONTAP 9.5

Configure NVE or NAE on controllers running ONTAP 9.5

Steps

1. Verify that the key management servers are still available, their status, and their authentication key information:

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

2. Add the key management servers listed in the previous step 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.

- Restore authentication for external key manager:

```
security key-manager external restore
```

This command needs the Onboard Key Manager (OKM) passphrase.

For more information, see the Knowledge Base article [How to restore external key manager server configuration from the ONTAP boot menu](#).

- Restore authentication for OKM:

```
security key-manager setup -node node_name
```

After you finish

Check if any volumes were taken offline because authentication keys were not available or External Key Management servers could not be reached. Bring those volumes back online by using the `volume online` command.

Decommission 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. Refer to [References](#) to link to the *NetApp Support Site* and log in.
2. Select **Products > My Products** from the menu.
3. On the **View Installed Systems** page, choose which **Selection Criteria** you want to use to display information about your system.

You can choose one of the following to locate your system:

- Serial Number (located on the back of the unit)
- Serial Numbers for My Location

4. Select **Go!**

A table displays cluster information, including the serial numbers.

5. Locate the cluster in the table and select **Decommission this system** from the Product Tool Set drop-down menu.

Resume SnapMirror operations

You can resume SnapMirror transfers that were quiesced before upgrade and resume the SnapMirror relationships. The updates are on schedule after the upgrade is completed.

Steps

1. Verify the SnapMirror status on the destination:

```
snapmirror show
```

2. Resume the SnapMirror relationship:

```
snapmirror resume -destination-vserver vs_server_name
```

Troubleshoot

Aggregate relocation failures

Aggregate relocation (ARL) might fail at different points during the upgrade.

Check for aggregate relocation failure

During the procedure, ARL might fail in Stage 2, Stage 3, or Stage 5.

Steps

1. Enter the following command and examine the output:

```
storage aggregate relocation show
```

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, refer to [References](#) to link to the *ONTAP 9 Commands: Manual Page Reference*.

Aggregates originally on node1 are owned by node4 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 correctly, 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 [Prepare the nodes for 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 [Prepare the nodes for upgrade](#) and note any aggregates that were not correctly relocated.
3. Relocate the aggregates left behind on node4:

```
storage aggregate relocation start -node node4 -aggr aggr_node_1 -destination node3
```

Do not use the `-ndo-controller-upgrade` parameter during this relocation.

4. 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 the pre-check phase

Node1 or node2 crashes before the pre-check phase with HA pair still enabled

If either node1 or node2 crashes before the pre-check phase, no aggregates have been relocated yet and the HA pair 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.

Reboots, panics, or power cycles during first resource-release phase

Node1 crashes during the first resource-release phase with HA pair still enabled

Some or all aggregates have been relocated from node1 to node2, and HA pair is still enabled. Node2 takes over node1's root volume and any non-root aggregates that were not relocated.

About this task

Ownership of aggregates that were relocated look the same as the ownership of non-root aggregates that were taken over because the home owner has not changed.

When node1 enters the `waiting for giveback` state, node2 gives back all of the node1 non- root aggregates.

Steps

1. After node1 is booted up, all the non-root aggregates of node1 have moved back to node1. You must perform a manual aggregate relocation of the aggregates from node1 to node2:

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

2. Continue with the node-pair upgrade procedure.

Node1 crashes during the first resource-release phase while HA pair is disabled

Node2 does 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.

Node2 fails during the first resource-release phase with HA pair still enabled

Node1 has relocated some or all of its aggregates to node2. The HA pair is enabled.

About this task

Node1 takes over all of node2's aggregates as well as any of its own aggregates that it had relocated to node2. When node2 boots up, the aggregate relocation is completed automatically.

Steps

1. Bring up node2.
2. Continue with the node-pair upgrade procedure.

Node2 crashes during the first resource-release phase and after HA pair is disabled

Node1 does not take over.

Steps

1. Bring up node2.

A client outage occurs for all aggregates while node2 is booting up.

2. Continue with the rest of the node-pair upgrade procedure.

Reboots, panics, or power cycles during the first verification phase

Node2 crashes during the first verification phase with HA pair disabled

Node3 does not take over following a node2 crash as the HA pair is already disabled.

Steps

1. Bring up node2.

A client outage occurs for all aggregates while node2 is booting up.

2. Continue with the node-pair upgrade procedure.

Node3 crashes during the first verification phase with HA pair disabled

Node2 does 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.

Reboots, panics, or power cycles during first resource-regain phase

Node2 crashes during the first resource-regain phase during aggregate relocation

Node2 has relocated some or all of its aggregates from node1 to node3. Node3 serves data from aggregates that were relocated. The HA pair is disabled and hence there is no takeover.

About this task

There is client outage for aggregates that were not relocated. On booting up node2, the aggregates of node1 are relocated to node3.

Steps

1. Bring up node2.
2. Continue with the node-pair upgrade procedure.

Node3 crashes during the first resource-regain phase during aggregate relocation

If node3 crashes while node2 is relocating aggregates to node3, the task continues after node3 boots up.

About this task

Node2 continues to serve remaining aggregates, but aggregates that were already relocated to node3 encounter client outage while node3 is booting up.

Steps

1. Bring up node3.
2. Continue with the controller upgrade.

Reboots, panics, or power cycles during post-check phase

Node2 or node3 crashes during the post-check phase

The HA pair is disabled hence this is no takeover. There is a client outage for aggregates belonging to the node that rebooted.

Steps

1. Bring up the node.
2. Continue with the node-pair upgrade procedure.

Reboots, panics, or power cycles during second resource-release phase

Node3 crashes during the second resource-release phase

If node3 crashes while node2 is relocating aggregates, the task continues after node3 boots up.

About this task

Node2 continues to serve remaining aggregates but aggregates that were already relocated to node3 and node3's own aggregates encounter client outages while node3 is booting.

Steps

1. Bring up node3.
2. Continue with the controller upgrade procedure.

Node2 crashes during the second resource-release phase

If node2 crashes during aggregate relocation, node2 is not taken over.

About this task

Node3 continues to serve the aggregates that have been relocated, but the aggregates owned by node2 encounter client outages.

Steps

1. Bring up node2.
2. Continue with the controller upgrade procedure.

Reboots, panics, or power cycles during the second verification phase

Node3 crashes during the second verification phase

If node3 crashes during this phase, takeover does not happen since HA is already disabled.

About this task

There is an outage for non-root aggregates that were already relocated until node3 reboots.

Steps

1. Bring up node3.

A client outage occurs for all aggregates while node3 is booting up.

2. Continue with the node-pair upgrade procedure.

Node4 crashes during the second verification phase

If node4 crashes during this phase, takeover does not happen. Node3 serves data from the aggregates.

About this task

There is an outage for non-root aggregates that were already relocated until node4 reboots.

Steps

1. Bring up node4.
2. Continue with the node-pair upgrade procedure.

Issues that can arise in multiple stages of the procedure

Some issues can occur during 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 Possible	State Description
node1	node2	false	Unknown
node2	node1	false	Node owns partner aggregates as part of the non-disruptive head upgrade procedure. Takeover is not possible: Storage failover is disabled.

The following example shows the output of the `storage failover show` command after a reboot or panic:

```
cluster::> storage failover show
```

Node	Partner	Takeover Possible	State Description
node1	node2	-	Unknown
node2	node1	false	Waiting for node1, Partial giveback, Takeover is not possible: Storage failover is disabled

Although the output says that a node is in partial giveback and that storage failover is disabled, you can disregard this message.

Steps

No action is required; continue with the node-pair upgrade procedure.

LIF migration failure

After you migrate LIFs, they might not come online after migration in Stage 2, Stage 3, or Stage 5.

Steps

1. Verify that the 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.

2. Check the physical connectivity of the network cable if the physical state of the port is down.

References

When performing the procedures in this content, you might need to consult reference content or go to reference websites.

Reference content

Content specific to this upgrade are listed in the table below.

Content	Description
Administration overview with the CLI	Describes how to administer ONTAP systems, shows you how to use the CLI interface, how to access the cluster, how to manage nodes, and much more.
Decide whether to use System Manager or the ONTAP CLI for cluster setup	Describes how to set up and configure ONTAP.
Disk and aggregate management with the CLI	Describes how to manage ONTAP physical storage using the CLI. It shows you how to create, expand, and manage aggregates, how to work with Flash Pool aggregates, how to manage disks, and how to manage RAID policies.
HA pair management	Describes how to install and manage high-availability clustered configurations, including storage failover and takeover/giveback.
Logical storage management with the CLI	Describes how to efficiently manage your logical storage resources, using volumes, FlexClone volumes, files, and LUNs, FlexCache volumes, deduplication, compression, qtrees, and quotas.
MetroCluster Management and Disaster Recovery	Describes how to perform MetroCluster switchover and switchback operations, both in planned maintenance operations, or in the event of a disaster.
MetroCluster Upgrade and Expansion	Provides procedures for upgrading controller and storage models in the MetroCluster configuration, transitioning from a MetroCluster FC to a MetroCluster IP configuration, and expanding the MetroCluster configuration by adding additional nodes.
Network Management	Describes how to configure and manage physical and virtual network ports (VLANs and interface groups), LIFs, routing, and host-resolution services in clusters; optimize network traffic by load balancing; and monitor the cluster by using SNMP.
ONTAP 9.0 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.0 commands.
ONTAP 9.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.1 commands.
ONTAP 9.2 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.2 commands.
ONTAP 9.3 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.3 commands.

Content	Description
ONTAP 9.4 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.4 commands.
ONTAP 9.5 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.5 commands.
ONTAP 9.6 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.6 commands.
ONTAP 9.7 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.7 commands.
ONTAP 9.8 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.8 commands.
ONTAP 9.9.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.9.1 commands.
ONTAP 9.10.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.10.1 commands.
SAN management with the CLI	Describes how to configure and manage LUNs, igroups, and targets using the iSCSI and FC protocols, and namespaces and subsystems using the NVMe/FC protocol.
SAN configuration reference	Contains information about FC and iSCSI topologies and wiring schemes.
Upgrade by moving volumes or storage	Describes how to quickly upgrade controller hardware in a cluster by moving storage or volumes. Also describes how to convert a supported model to a disk shelf.
Upgrade ONTAP	Contains instructions for downloading and upgrading ONTAP.
Use "system controller replace" commands to upgrade controller hardware introduced in ONTAP 9.15.1 and later	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers introduced in ONTAP 9.15.1 and later by using "system controller replace" commands.
Use "system controller replace" commands to upgrade controller models in the same chassis	Describes the aggregate relocation procedures needed to non-disruptively upgrade a system, keeping the old system chassis and disks.
Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers running ONTAP 9.8 by using "system controller replace" commands.
Use aggregate relocation to manually upgrade controller hardware running ONTAP 9.8 or later	Describes the aggregate relocation procedures needed to perform manual non-disruptive controller upgrades running ONTAP 9.8 or later.
Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.5 to ONTAP 9.7	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers running ONTAP 9.5 to ONTAP 9.7 by using "system controller replace" commands.
Use aggregate relocation to manually upgrade controller hardware running ONTAP 9.7 or earlier	Describes the aggregate relocation procedures needed to perform manual non-disruptive controller upgrades running ONTAP 9.7 or earlier.

Reference sites

The [NetApp Support Site](#) also contains documentation about network interface cards (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.

Access [ONTAP 9 documentation](#).

Access the [Active IQ Config Advisor](#) tool.

Manually upgrade controller hardware running ONTAP 9.8 or later

Learn about this ARL upgrade procedure

This procedure describes how to upgrade controller hardware using manual aggregate relocation (ARL) on systems running ONTAP 9.8 or later.

You can use this ARL procedure if you are performing one of the following upgrades:

- FAS system to FAS system
- AFF system to AFF system

You can only upgrade to a replacement system in the same series:

- AFF A-Series system to AFF A-Series system
- AFF C-Series system to AFF C-Series system
- ASA system to ASA system

ASA upgrades to an ASA r2 replacement system aren't supported. For information on migrating data from ASA to ASA r2, see [Enable data access from SAN hosts to your ASA r2 storage system](#).

You can only upgrade to a replacement system in the same series:

- ASAA-Series system to ASAA-Series system
- ASA C-Series system to ASA C-Series system

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 confirm 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.



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*.

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.

Important information:

- This procedure is complex and assumes that you have advanced ONTAP administration skills. You also must read and understand the [guidelines for upgrading controllers with ARL](#) and the [ARL upgrade workflow](#) before beginning the upgrade.
- 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.
- 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 high-availability (HA) pair in the cluster.
- In addition to non-MetroCluster configurations, this procedure applies to Fabric MetroCluster four-node and eight-node configurations running ONTAP 9.8 and later.
 - For MetroCluster configurations running ONTAP 9.7 and earlier, go to [References](#) to link to *Using Aggregate Relocation to Manually Upgrade Controller Hardware Running ONTAP 9.7 or Earlier*.
 - For MetroCluster IP configurations and additional upgrade options for Fabric MetroCluster configurations, go to [References](#) to link to the *MetroCluster Upgrade and Expansion* content.

Decide whether to use this aggregate relocation procedure

This procedure describes how to upgrade controller hardware using manual aggregate relocation (ARL) on systems running ONTAP 9.8 or later. You should only use this complex procedure if you're an experienced ONTAP administrator.

Use this content under the following circumstances:

- You don't 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.
- Your controllers are running ONTAP 9.8 or later.
- You have a system that uses Fabric MetroCluster 4-node and 8-node configurations running ONTAP 9.8 or later.
- You have hybrid aggregates on your system.



- If you're upgrading a system by swapping controller modules in the same chassis, such as AFF A800 or AFF C800, NetApp strongly recommends using the upgrade procedure that [upgrades controller models using ARL, keeping the existing system chassis and disks](#). This ARL procedure includes the steps that ensure the internal disks remain secure in the chassis when you remove and install the controllers during the upgrade procedure.

[Learn about the supported system upgrade combinations using ARL, keeping the existing system chassis and disks.](#)

- You can use NetApp Storage Encryption (NSE), NetApp Volume Encryption (NVE), and NetApp Aggregate Encryption (NAE) with this procedure.

If you prefer a different method of upgrading the controller hardware and are willing to do volume moves, refer to [References](#) to link to *Upgrade by moving volumes or storage*.

Refer to [References](#) to link to the *ONTAP 9 Documentation Center* where you can access ONTAP 9 product documentation.

Choose a different hardware upgrade procedure

- [Review the alternative ARL methods available for upgrading controller hardware.](#)
- If you prefer a different method of upgrading the controller hardware and are willing to do volume moves, refer to [References](#) to link to *Upgrade by moving volumes or storage*.

Related information

Refer to [References](#) to link to the *ONTAP 9 Documentation*.

ARL upgrade workflow

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

Upgrade the node pair


To upgrade the node pair, you must 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.

ARL upgrade sequence overview

During the procedure, you upgrade the original controller hardware with the replacement controller hardware, one controller at a time, taking advantage of the HA pair 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:

Stage	Steps
Stage 1: Prepare for upgrade	<p>During Stage 1, if required, you confirm that internal disk drives do not contain root aggregates or data aggregates, prepare the nodes for the upgrade, and run a series of prechecks. If required, you rekey disks for Storage Encryption and prepare to netboot the new controllers.</p> <p>Aggregate ownership at the end of Stage 1:</p> <ul style="list-style-type: none"> Node1 is the home owner and current owner of the node1 aggregates. Node2 is the home owner and current owner of the node2 aggregates.

Stage	Steps
Stage 2: Retire node1	<p>During Stage 2, you relocate non-root aggregates from node1 to node2 and move non-SAN data LIFs owned by node1 to node2, including failed or vetoed aggregates. You also record the necessary node1 information for use later in the procedure and retire node1.</p> <p>Aggregate ownership at the end of Stage 2:</p> <ul style="list-style-type: none"> • Node1 is the home owner of node1 aggregates. • Node2 is the current owner of node1 aggregates. • Node2 is the home owner and current owner of node2 aggregates.
Stage 3: Install and boot node3	<p>During Stage 3, you install and boot node3, map the cluster and node-management ports from node1 to node3, 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.</p> <p>Aggregate ownership at the end of Stage 3:</p> <ul style="list-style-type: none"> • Node2 is the home owner of node2 aggregates but not the current owner. • Node3 is the home owner and current owner of aggregates originally belonging to node1. • Node2 is the home owner and current owner of aggregates belonging to node2 but not the home owner.
Stage 4: Retire node2	<p>During Stage 4, you record the necessary node2 information for use later in the procedure and then retire node2.</p> <p>No changes occur in aggregate ownership.</p>
Stage 5: Install and boot node4	<p>During Stage 5, you install and boot node4, map the cluster and node-management ports from node2 to node4, and move data LIFs and SAN LIFs belonging to node2 from node3 to node4. You also relocate node2 aggregates from node3 to node4 and move the data LIFs and SAN LIFs owned by node2 to node3.</p> <p>Aggregate ownership at the end of Stage 5:</p> <ul style="list-style-type: none"> • Node3 is the home owner and current owner of the aggregates that originally belonged to node1. • Node4 is the home owner and current owner of aggregates that originally belonged to node2.

Stage	Steps
Stage 6: Complete the upgrade	<p>During Stage 6, you confirm that the new nodes are set up correctly and set up Storage Encryption or NetApp Volume Encryption if the new nodes are encryption-enabled. You should also decommission the old nodes resume SnapMirror operations.</p> <div>  <p>The storage virtual machine (SVM) disaster recovery updates will not be interrupted as per the schedules assigned.</p> </div> <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 9.8 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 9.8 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.
 - An AFF system to an AFF system.
- For some ARL controller upgrades you can use temporary cluster ports on the replacement controller for the upgrade. For example, if you upgrade from an AFF A300 to an AFF A400 system, depending on the AFF A400 configuration, you can use any of the two mezzanine ports or add a four-port 10GbE network interface card to provide temporary cluster ports. After you complete a controller upgrade using temporary cluster ports, you can nondisruptively migrate clusters to 100GbE ports on the replacement controller.
- Controller upgrade using ARL is supported on systems configured with SnapLock Enterprise and SnapLock Compliance volumes.

You must verify whether the ARL can be performed on the original and replacement controllers. You must 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. To access this information, refer to [References](#) to link to the *Hardware Universe*. 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 must 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. For more information about cluster mixing rules, refer to [References](#) to link to the *Hardware Universe*.



Both systems are either high-availability (HA) or non-HA. 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, contact technical support.



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. Refer to [References](#) to link to the *Hardware Universe* on the NetApp Support Site for details about specific platforms.

If you have a system with more than two cluster ports per node, such as an FAS8080 or an AFF8080 system, before you start the upgrade, you must migrate and re-home the cluster LIFs to two cluster ports per node. If you perform the controller upgrade with more than two cluster ports per node, cluster LIFs might be missing on the new controller after the upgrade.

Upgrades not supported for ARL

You cannot perform the following upgrades:

- To or from controllers that cannot run ONTAP 9.8 or later.
- To replacement controllers that do not support the disk shelves connected to the original controllers.

For disk-support information, refer to [References](#) to link to the *Hardware Universe*.

- 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, refer to [References](#) to link to *Upgrade by moving volumes or storage* and go to the procedure *Upgrading a pair of nodes running clustered Data ONTAP by moving volumes*.



If you want to upgrade ONTAP on nodes in a cluster, refer to [References](#) to link to *Upgrade ONTAP*.

Assumptions and terminology

This document is written with the following assumptions:

- 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 because this procedure assumes that the replacement controller hardware is new and has not been used. You must contact technical support if the replacement controller hardware was previously used, especially if the controllers were running Data ONTAP in 7-Mode.

- You read and understand the guidelines for upgrading the pair of nodes.



Do not try to clear the NVRAM contents. If you 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 partner LIFs for each storage virtual machine (SVM), on the HA pair. If you do not have local and partner 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 9.8 or Later

Some features require 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 and will continue to work. However, using unlicensed features on the controller might put you out of compliance with your license agreement, so you must install the new license key or keys for the new controller after the upgrade is complete.

All license keys are 28 uppercase alphabetic characters in length. Refer to [References](#) to link to the *NetApp Support Site* where you can obtain new 28-character license keys for ONTAP 9.8. or later. The keys are available in the *My Support* section under *Software licenses*. If the site does not have the license keys you need, contact your NetApp sales representative.

For detailed information about licensing, go to [References](#) to link to the *System Administration Reference*.

Storage Encryption

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

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

Troubleshoot

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.

If any problems occur while upgrading the controllers, refer to the [Troubleshoot](#) section. The information about failures that can occur is listed by the phase of the procedure in the [ARL upgrade sequence](#).

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

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 must 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

Go to the [References](#) section to access the list of reference documents required for this upgrade.

Worksheet: Information to collect before and during controller upgrade

You should gather certain information to support upgrading the original nodes. This information includes node IDs, port and LIF details, licensing keys, and IP addresses.

You can use the following worksheet to record the 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>Prepare the nodes for the upgrade</i>	Stage 3: <i>Install and boot node3</i> Stage 5: <i>Install and boot node4</i> Stage 6: <i>Decommission 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>Prepare the nodes for the upgrade</i>	Throughout the procedure to verify that aggregates and volumes remain online except during brief relocation	
Output of commands network port vlan show and network port ifgrp show	Stage 1: <i>Prepare the nodes for the upgrade</i>	Stage 3: <i>Map ports from node1 to node3</i> Stage 5: <i>Map ports from node2 to node4</i>	
(SAN environments only) Default configuration of FC ports	Stage 1: <i>Prepare the nodes for the upgrade</i>	When configuring FC ports on the new nodes	

Information needed	When collected	When used	Collected Information
IP address of SPs	Stage 1: <i>Prepare the nodes for the upgrade</i>	Stage 6: <i>Confirm that the new controllers are set up correctly</i>	
License keys	Stage 1: <i>Prepare the nodes for the upgrade</i>	Stage 6: <i>Confirm that the new controllers are set up correctly</i>	
IP address for the External Key Management server	Stage 1: <i>Rekey disks for Storage Encryption</i>	Stage 6: <i>Set 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>Prepare to netboot</i>	Stage 3: <i>Install and boot node3</i> Stage 5: <i>Install and boot node4</i>	
Non-SAN data LIFs owned by node1	Stage 2: <i>Move nonSAN data LIFs owned by node1 to node2</i>	Later in the section	
Cluster, intercluster, node-management, cluster-management, and physical ports	Stage 2: <i>Record node1 information</i>	Stage 3: <i>Install and boot node3</i> Stage 3: <i>Map ports from node1 to node3</i>	
Ports on new nodes	Stage 3: <i>Map ports from node1 to node3</i>	Later in the section and in the section <i>Map ports from node2 to node4</i>	
Available ports and broadcast domains on node3	Stage 3: <i>Map 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>Move nonSAN data LIFs owned by node2 to node3</i>	Later in the section	
Cluster, intercluster, node-management, cluster-management, and physical ports	Stage 4: <i>Record node2 information</i>	Stage 5: <i>Install and booting node4</i> Stage 5: <i>Map ports from node2 to node4</i>	
Cluster network ports on node4	Stage 5: <i>Map ports from node2 to node4</i>	Later in the section	

Information needed	When collected	When used	Collected Information
Available ports and broadcast domains on node4	Stage 5: <i>Map 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>Set up Storage Encryption on the new nodes</i>	Later in the section	

Stage 1. Prepare for upgrade

Determine 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 confirm 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 [Prepare the nodes for 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
```

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)
...

```

3. 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.

4. Enter the following command on both controllers:

```
aggr status -r
```

The system displays the aggregate status of the node, as shown in the partial output in 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
...
```



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

5. 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 the previous step, "aggr2" uses internal drives, as indicated by the shelf ID of "0".

6. Enter the following command on both controllers:

```
aggr status -v
```

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,
           64-bit           raidsize=14, 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
           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.

- 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.
- Exit the nodeshell by entering the following command on each controller:

```
exit
```

- Take one of the following actions:

If the controllers....	Then...
Do not contain any aggregates on internal disk drives	Continue with this procedure.

If the controllers....	Then...
Contain aggregates but no volumes on the internal disk drives	<p>Continue with this procedure.</p> <div>  <p>Before you continue, you must place the aggregates offline, and then destroy the aggregates on the internal disk drives. Refer to References to link to the <i>Disk and aggregate management with the CLI</i> content for information about managing aggregates.</p> </div>
Contain non-root volumes on the internal drives	<p>Continue with this procedure.</p> <div>  <p>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. Refer to References to link to the <i>Disk and aggregate management with the CLI</i> content for information about moving volumes.</p> </div>
Contain root volumes on the internal drives	<p>Do not continue with this procedure.</p> <p>You can upgrade the controllers by referring to References to link to the <i>NetApp Support Site</i> and using the procedure <i>Upgrading the controller hardware on a pair of nodes running clustered Data ONTAP by moving volumes</i>.</p>
Contain non-root volumes on the internal drives and you cannot move the volumes to external storage	<p>Do not continue with this procedure.</p> <p>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>. Refer to References to link to the <i>NetApp Support Site</i> where you can access this procedure.</p>

Prepare the nodes for upgrade

Before you can replace the original nodes, you must confirm 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 must collect information about the original nodes and, if the cluster is in a SAN environment, confirm that all the nodes in the cluster are in quorum.

Steps

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

Refer to [References](#) to link to *HA pair management* and follow the *Best practices for HA pairs* section. 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.



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

For information about unlocking the user accounts, refer to [References](#) to link to the *System Administration Reference*.

- b. Refer to [References](#) to link to the *NetApp Support Site* and download the Performance and Statistics Collector (Perfstat Converged).

The 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. Refer to [References](#) to link to the *NetApp Support Site* and open a support case on the NetApp Support Site.

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 must physically check node3 and node4 to see if the NVMEM or NVRAM batteries are charged. For information about the LEDs for the model of node3 and node4, refer to [References](#) to link to the *Hardware Universe*.



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 must netboot the new controllers after you install them. For instructions on how to upgrade ONTAP, refer to [References](#) to link 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 node2:

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

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

```
cluster::> network interface show -role cluster
```

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

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. Confirm 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
```

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. For detailed information about the commands, refer to [References](#) to link to *HA pair management*.

9. Confirm that neither node1 nor node2 owns the 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
```

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, homename, 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 .

If the command in Step 9...	Then...
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
```

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. Confirm 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
```

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, refer to [References](#) to link to *Disk and aggregate management with the CLI*, *Logical storage management with the CLI*, and *HA pair management* to configure storage for the HA pair.

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

```
cluster show
```

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. Set the privilege level to advanced:

```
set -privilege advanced
```

15. Confirm that node1 and node2 are running the same ONTAP release:

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

The following example shows the output of the command:

```
cluster::*> system node image show -node node1,node2 -iscurrent true
```

Node	Image	Is Default	Is Current	Version	Install 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.

16. 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
```

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
```

Current Is	Logical	Status	Network	Current	
Vserver	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.

17. 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 *
```

For detailed information about the `network interface revert` command, refer to [References](#) to link to the *ONTAP 9 Commands: Manual Page Reference*.

18. 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 *Disk and aggregate management with the CLI*. (Refer to [References](#) to link to *Disk and aggregate management with the CLI*.)

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



- You will use this information later in the procedure.
- If you have a system with more than two cluster ports per node, such as an FAS8080 or an AFF8080 system, before you start the upgrade, you must migrate and re-home the cluster LIFs to two cluster ports per node. If you perform the controller upgrade with more than two cluster ports per node, cluster LIFs might be missing on the new controller after the upgrade.

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

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



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
```



You can use the information to identify parts or accessories that you might want to transfer to node3 or node4.

- 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
```



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
```



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.

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

```
network port ifgrp show
```

```
network port vlan 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.

21. Complete the following substeps on both node1 and node2 to confirm 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.

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


```
cluster::> network interface failover-groups show
```

Vserver	Group	Targets
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.

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

```
cluster::> network port vlan show
```

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.
22. Take one of the following actions:

If interface groups or VLANs are...	Then...
On node1 or node2	Complete Step 23 and Step 24 .
Not on node1 or node2	Go to Step 24 .

23. 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 vservice_name -data-protocol iscsi|fc
```

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 Vserver8970 -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 `network interface show` command with the `-data-protocol nfs|cifs` parameters.

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 vs1 -data-protocol iscsi|fc
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
vs1	vs1_lif1	up/down	172.17.176.20/24	node1	0d	true

24. 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
```

The system displays a message similar to the following example:

```
cluster::*> cluster kernel-service show
```

Master Node	Cluster Node	Quorum Status	Availability Status	Operational 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, refer to [References](#) to link to the *System Administration Reference*.

d. Return to the administrative privilege level:

```
set -privilege admin
```

25. Take one of the following actions:

If the cluster...	Then...
Has SAN configured	Go to Step 26 .
Does not have SAN configured	Go to Step 29 .

26. 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|fc -home-node node_name
```

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 iscsi|fc
          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 following command:

```
network interface show -instance -data-protocol iscsi|fc
```

27. 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
```

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.

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. Confirm that the SP status is online.
c. Confirm 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, refer to [References](#) to link to 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 *
```

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 at the *NetApp Support Site*. Refer to [References](#) to link to *NetApp Support Site*.

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
```

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...	Go to Step 34 .

If the original system...	Then...
Does not have AutoSupport enabled...	<p>Enable AutoSupport by following the instructions in the <i>System Administration Reference</i>. (Refer to References to link to the <i>System Administration Reference</i>.)</p> <p>Note: 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.</p>

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, refer to [References](#) to link to 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"
```



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.

36. If your system uses self-encrypting drives, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use on the HA pair that you are upgrading. ONTAP software supports two types of self-encrypting drives:

- FIPS-certified NetApp Storage Encryption (NSE) SAS or NVMe drives
- Non-FIPS self-encrypting NVMe drives (SED)



You cannot mix FIPS drives with other types of drives on the same node or HA pair.

You can mix SEDs with non-encrypting drives on the same node or HA pair.

[Learn more about supported self-encrypting drives.](#)

Manage authentication keys using the Onboard Key Manager

You can use the Onboard Key Manager (OKM) to manage authentication keys. If you have the OKM set up, you must record the passphrase and backup material before beginning the upgrade.

Steps

1. Record the cluster-wide passphrase.

This is the passphrase that was entered when the OKM was configured or updated using the CLI or REST API.

2. Back up the key-manager information by running the `security key-manager onboard show-backup` command.

Quiesce the SnapMirror relationships

Before you netboot the system, you must confirm that all the SnapMirror relationships are quiesced. When a SnapMirror relationship is quiesced, it remains quiesced across reboots and failovers.

Steps

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

```
snapmirror show
```



If the status is `Transferring`, you must abort those transfers:

```
snapmirror abort -destination-vserver vservers name
```

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

2. Quiesce all relationships between the cluster:

```
snapmirror quiesce -destination-vserver *
```

Prepare for netboot

After you physically rack node3 and node4 later in the procedure, you might need to netboot them. The term *netboot* means you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you must put a copy of the ONTAP 9 boot image onto a web server that the system can access.

Before you begin

- Verify that you can access a HTTP server with the system.
- Refer to [References](#) to link to the *NetApp Support Site* and download the necessary system files for your platform and the correct version of ONTAP.

About this task


You must netboot the new controllers if they do not have the same 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 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 [Stage 3: Install and boot node3](#).

Steps

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

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

You will use information in the directories in [Stage 3](#).

Stage 2. Relocate and retire node1

Relocate non-root aggregates from node1 to node2

Before you can replace node1 with node3, you must 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. Set the privilege level to advanced:

```
set -privilege advanced
```

b. Enter the following command:

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

c. 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.

d. 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.



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:

If relocation...	Then..
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. Relocate any failed or vetoed aggregates:</p> <pre>storage aggregate relocation start -node node1 - destination node2 -aggregate-list * -ndo -controller-upgrade true</pre> <p>d. When prompted, enter y.</p> <p>e. Return to the admin level:</p> <pre>set -privilege admin</pre> <p>If necessary, you can force the relocation using one of the following methods:</p> <ul style="list-style-type: none"> ◦ Override veto checks: <pre>storage aggregate relocation start -override -vetoes true -ndo-controller-upgrade</pre> ◦ Override destination checks: <pre>storage aggregate relocation start -override -destination-checks true -ndo-controller -upgrade</pre> <p>Refer to References to link to the <i>Disk and aggregate management with the CLI</i> content 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:

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

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,

normal
aggr_2      825.0GB 825.0GB      0% online      1 node2
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
```

5. 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 vserver-name -volume volume-name
```

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

6. 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.
```

7. Verify that node1 does not own any non-root aggregates that are online:

```
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.

Move NAS data LIFs owned by node1 to node2

Before you can replace node1 with node3, you must move the NAS 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

1. List all the NAS data LIFs hosted on node1 by entering the following command and capturing the output:

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

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

Is	Logical	Status	Network	Current	Current
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					
	rads2	up/down	10.63.0.52/24	node1	e1b
true					
vs1					
	lif1	up/up	192.17.176.120/24	node1	e0c
true					
	lif2	up/up	172.17.176.121/24	node1	e1a
true					

2. Modify the auto revert settings of all the LIFs on node1 and node2:

```
network interface modify -vserver Vserver_name -lif LIF_name -auto-revert
false
```

3. Take the following steps to migrate any NAS 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 Vserver_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 Vserver_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 Vserver-name -lif LIF_name -home-node node_to_upgrade -home-port netport ifgrp -status -admin down</pre>

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

```
network interface migrate -vserver Vserver-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. Modify the home node of the migrated LIFs:

```
network interface modify -vserver Vserver-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 Vserver-name -lif LIF_name
-destination-node node_name -destination-port port_name
```

```
network interface modify -vserver Vserver-name -lif LIF_name -home-node
node_name -home-port port_name
```

10. 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 -home-node
nodename -status-admin up
```



For MetroCluster configurations, 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 storage virtual machine (SVM). Enter the following command from the corresponding source SVM on the remote site to reallocate the destination LIF to an appropriate port:

```
metrocluster vsync resync -vserver Vserver_name
```

11. 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
```

Record node1 information

Before you can shut down and retire node1, you must 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
```

The system displays output similar to the following example:

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

Vserver	Destination	Gateway	Metric
-----	-----	-----	-----
iscsi vsync	0.0.0.0/0	10.10.50.1	20
node1	0.0.0.0/0	10.10.20.1	10
....			
node2	0.0.0.0/0	192.169.1.1	20

2. Enter the following command and capture its output:

```
vsync services name-service dns show
```

The system displays output similar to the following example:

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

Vserver	State	Domains	Name Servers
node 1 2 10.10.60.10, 10.10.60.20	enabled	alpha.beta.gamma.netapp.com	
vs_base1 10.10.60.10, 10.10.60.20	enabled	alpha.beta.gamma.netapp.com, beta.gamma.netapp.com,	
...			
...			
vs_peer1 10.10.60.10, 10.10.60.20	enabled	alpha.beta.gamma.netapp.com, gamma.netapp.com	

- 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
```

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
```

Current Is	Logical	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	
vserver1	cluster mgmt	up/up	192.168.x.xxx/24	node1	e0c
true					
node1	intercluster	up/up	192.168.x.xxx/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.x.xxx/24	node1	e0c
true					

5 entries were displayed.



Your system might not have intercluster LIFs.

- Capture the information in the output of the command in [Step 3](#) to use in the section [Map ports from node1 to node3](#).

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

- Enter the following command on node1:

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

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

```
sti8080mcc-htp-008::> network port show -node sti8080mcc-htp-008 -type
physical
```

Node: sti8080mcc-htp-008

Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Oper	Health Status	Ignore Health Status
----	-----	-----	----	----	-----	-----	
e0M	Default	Mgmt	up	1500	auto/1000	healthy	false
e0a	Default	Default	up	9000	auto/10000	healthy	false
e0b	Default	-	up	9000	auto/10000	healthy	false
e0c	Default	-	down	9000	auto/-	-	false
e0d	Default	-	down	9000	auto/-	-	false
e0e	Cluster	Cluster	up	9000	auto/10000	healthy	false
e0f	Default	-	up	9000	auto/10000	healthy	false
e0g	Cluster	Cluster	up	9000	auto/10000	healthy	false
e0h	Default	Default	up	9000	auto/10000	healthy	false

9 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
```

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 [Map ports from node1 to node3](#).

10. Take one of the following actions:

If you...	Then...
Recorded the NVRAM System ID number in the section Prepare the nodes for the upgrade .	Go on to the next section, Retire node1 .
Did not record the NVRAM System ID number in the section Prepare the nodes for the upgrade	Complete Step 11 and Step 12 and then continue to Retire node1 .

11. Enter the following command on either controller:

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

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 [Install and boot node3](#).

Retire node1

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

Steps

1. Verify the number of nodes in the cluster:

```
cluster show
```

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



If you do not disable storage failover, a controller upgrade failure can occur which can disrupt data access and lead to data loss.

3. Verify that storage failover was disabled:

```
storage failover show
```

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 [Troubleshoot](#) section.

5. Take one of the following actions:

If the cluster is...	Then...
A two-node cluster	Go to Step 6 .
A cluster with more than two nodes	Go to Step 8 .

6. Access the advanced privilege level on either node:

```
set -privilege advanced
```

7. 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](#).

8. 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. Refer to [References](#) to link to the *System Administration Reference* for more information.



If you have a four-node cluster, epsilon might be on a node in a different HA pair in the cluster.

If you are upgrading an HA pair in a cluster with multiple HA pairs, you must move epsilon to the node of an HA pair that isn't undergoing a controller upgrade. For example, if you are upgrading nodeA/nodeB in a cluster with the HA pair configuration nodeA/nodeB and nodeC/nodeD, you must move epsilon to nodeC or nodeD.

The following example shows that node1 holds epsilon:

```
cluster::*> cluster show
```

Node	Health	Eligibility	Epsilon
node1	true	true	true
node2	true	true	false

9. 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
```

10. Transfer epsilon to node2 by marking epsilon true on node2:

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

11. Verify that the change to node2 occurred:

```
cluster show
```

```
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.

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

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

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

Enable Switchless Cluster: false/true
```

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

13. Return to the admin level:

```
set -privilege admin
```

14. Halt node1 from the node1 prompt:

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



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.

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

The node stops at the boot environment prompt.

16. 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 [Decommission the old system](#).

Stage 3. Install and boot node3

Install and boot node3

You must install node3 in the rack, transfer node1's connections to node3, boot node3, and install ONTAP. You must also 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 must netboot node3 if it doesn't have the same version of ONTAP 9 that is installed on node1. After you install node3, boot it 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. See [Prepare for netboot](#).

However, you don't need to netboot node3 if it has the same or a later version of ONTAP 9 that is installed on node1.



- For an AFF A800 or AFF C800 controller upgrade, you must ensure that all drives in the chassis are firmly seated against the midplane before removing node1. For more information, see [Replace the AFF A800 or AFF C800 controller modules](#).
- If you are upgrading a system with storage disks, you must complete this entire section and then go to [Configure FC ports on node3](#) and [Check and configure UTA/UTA2 ports on node3](#), 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 must put node3 into its own rack space, preferably close to the location of node1.

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



If you're upgrading to a system with both nodes in the same chassis, install node4 and node3 in the chassis. If you don't install both nodes in the same chassis, when you boot node3, it behaves as if it were in a dual-chassis configuration, and when you boot node4, the interconnect between the nodes doesn't come up.

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

The following references help you make proper cable connections. Go to [References](#) to link to them.

- *Installation and Setup Instructions* for the node3 platform
- The appropriate disk shelf procedure
- The *HA pair management* documentation

Cable the following connections:

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



You might 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 must move the FC-VI cable connections from node1 to node3. If the new host does not have an FC-VI card, you might need to move the FC-VI card.

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 later.



When you boot node3, you might see the following warning 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:

- 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 complete.




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

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

```
halt
```

The system stops at the boot environment prompt.

7. 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 8 .
Single-chassis configuration (with controllers in the same chassis)	<ol style="list-style-type: none">Switch the console cable from node3 to node4.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. <div> Leave node4 at the boot environment prompt; you will return to node4 in Install and boot node4.</div>If you see the warning message displayed in Step 4, follow the instructions in Step 5Switch the console cable back from node4 to node3.Go to Step 8.

8. Configure node3 for ONTAP:

```
set-defaults
```

9. If you have NetApp Storage Encryption (NSE) drives installed, perform the following steps:



If you have not already done so earlier in the procedure, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use.

- a. Set `bootarg.storageencryption.support` to `true` or `false`:

If the following drives are in use...	Then...
NSE drives that conform to FIPS 140-2 Level 2 self-encryption requirements	<code>setenv bootarg.storageencryption.support true</code>
NetApp non-FIPS SEDs	<code>setenv bootarg.storageencryption.support false</code>



You cannot mix FIPS drives with other types of drives on the same node or HA pair.

You can mix SEDs with non-encrypting drives on the same node or HA pair.

- b. Go to the special boot menu and select option (10) `Set Onboard Key Manager recovery secrets`.

Enter the passphrase and the backup information that you recorded earlier procedure. See [Manage authentication keys using the Onboard Key Manager](#).

10. If the version of ONTAP installed on node3 is the same or later than the version of ONTAP 9 installed on node1, list and reassign disks to the new node3:

```
boot_ontap
```



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 might result in service outages or data loss. Contact technical support if the replacement controller was previously used, especially if the controllers were running ONTAP running in 7-Mode.

11. Press CTRL-C to display the boot menu.


12. Take one of the following actions:

If the system you are upgrading...	Then...
Does <i>not</i> have the correct or current ONTAP version on node3	Go to Step 13 .
Has the correct or current version of ONTAP on node3	Go to Step 18 .

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



You must use the management port and IP as the netboot connection. Do not use a data LIF IP or else 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: <code>ifconfig e0M -auto</code>
Not running	<p>Manually configure the connection by entering the following command at the boot environment prompt: <code>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></code></p> <p><i>filer_addr</i> is the IP address of the storage system (mandatory). <i>netmask</i> is the network mask of the storage system (mandatory). <i>gateway</i> is the gateway for the storage system (mandatory). <i>dns_addr</i> is the IP address of a name server on your network (optional). <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> <div>  <p>Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p> </div>

14. Perform netboot on node3:

For...	Then...
FAS/AFF8000 series systems	<code>netboot</code> <code>http://<web_server_ip>/<path_to_webaccessible_directory>/netboot/kernel</code>
All other systems	<code>netboot</code> <code>http://<web_server_ip>/<path_to_webaccessible_directory>/<ontap_version>_image.tgz</code>

The `<path_to_the_web-accessible_directory>` leads to where you downloaded the `<ontap_version>_image.tgz` in [Step 1](#) in the section *Prepare for netboot*.



Do not interrupt the boot.

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

This menu option downloads and installs the new 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 ONTAP, and not upgrades of controllers.



Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the wrong image might install. This issue applies to all releases of ONTAP. The netboot procedure combined with option (7) *Install new software* wipes the boot media and places the same ONTAP version ONTAP on both image partitions.

16. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the following URL:

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

17. 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}
```

- 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.

18. Select **(5) Maintenance mode boot** by entering `5`, and then enter `y` when prompted to continue with the boot.
19. Before continuing, go to [Set the FC or UTA/UTA2 configuration on node3](#) to make any necessary changes to the FC or UTA/UTA2 ports on the node.

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

20. Find the system ID of node3:

```
disk show -a
```

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  HOME      DR
HOME                                NUMBER
-----
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)
.....
```



You might 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.

21. Reassign node1's spare disks, any disks belonging to the root, and any non-root aggregates that were not relocated to node2 earlier in [Relocate non-root aggregates from node1 to node2](#).

Enter the appropriate form of the `disk reassign` command based on whether your system has shared disks:



If you have shared disks, hybrid aggregates, or both on your system, you must use the correct `disk reassign` command from the following table.

If disk type is...	Then run the command...
With shared disks	<code>disk reassign -s <i>node1_sysid</i> -d <i>node3_sysid</i> -p <i>node2_sysid</i></code>
Without shared disks	<code>disk reassign -s <i>node1_sysid</i> -d <i>node3_sysid</i></code>

For the `node1_sysid` value, use the information captured in [Record node1 information](#). To obtain the value for `node3_sysid`, use the `sysconfig` command.



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)?
```

22. 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)?
```

23. 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)?
```

24. Enter *y*.

25. If you are upgrading from a system with external disks to a system that supports internal and external disks (AFF A800 systems, for example), set the node1 aggregate as root to confirm that node3 boots from the root aggregate of node1.



Warning: You must perform the following substeps in the exact order shown; failure to do so might cause an outage or even data loss.

The following procedure sets node3 to boot from the root aggregate of node1:

a. Check the RAID, plex, and checksum information for the node1 aggregate:

```
aggr status -r
```

b. Check the status of the node1 aggregate:

```
aggr status
```

c. Bring the node1 aggregate online, if necessary:

```
aggr_online root_aggr_from_node1
```

d. Prevent the node3 from booting from its original root aggregate:

```
aggr offline root_aggr_on_node3
```

- e. Set the node1 root aggregate as the new root aggregate for node3:

```
aggr options aggr_from_node1 root
```

- f. Verify that the root aggregate of node3 is offline and the root aggregate for the disks brought over from node1 is online and set to root:

```
aggr status
```



Failing to perform the previous substep might cause node3 to boot from the internal root aggregate, or it might cause the system to assume a new cluster configuration exists or prompt you to identify one.

The following shows an example of the command output:

```
-----  
      Aggr State      Status      Options  
aggr0_nst_fas8080_15 online  raid_dp, aggr  root, nosnap=on  
                        fast zeroed  
                        64-bit  
  
      aggr0 offline      raid_dp, aggr  diskroot  
                        fast zeroed  
                        64-bit  
-----
```

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

```
ha-config show
```

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 programmable ROM (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
```

```
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
```

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

27. 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?
```

28. Enter `y` at the prompt to confirm that you want to destroy the local mailboxes.

29. Exit maintenance mode:

```
halt
```

The system stops at the boot environment prompt.

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

```
date
```

31. On node3, check the date at the boot environment prompt:

```
show date
```

32. If necessary, set the date on node3:

```
set date mm/dd/yyyy
```

33. On node3, check the time at the boot environment prompt:

```
show time
```

34. If necessary, set the time on node3:

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

35. Verify the partner system ID is set correctly as noted in [Step 21](#) under `-p` switch:

```
printenv partner-sysid
```

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

```
setenv partner-sysid node2_sysid
```


Save the settings:

```
saveenv
```

37. Access the boot menu at the boot environment prompt:

```
boot_ontap menu
```

38. 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?:
```

39. Enter `y` at the prompt.

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



The system might reboot twice before displaying the mismatch warning.

40. Confirm the mismatch as shown in the following 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.

41. Log in to node3.

Set the FC or UTA/UTA2 configuration on node3

If node3 has onboard FC ports, onboard unified target adapter (UTA/UTA2) ports, or a UTA/UTA2 card, you must configure the settings before completing the rest of the procedure.

About this task

You might need to complete [Configure FC ports on node3](#) or [Check and configure UTA/UTA2 ports on node3](#), or both sections.



NetApp marketing materials might use the term "UTA2" to refer to CNA adapters and ports. However, the CLI uses the term "CNA".

If node3 does not have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card (for example, AFF and FAS systems introduced beginning with ONTAP 9.15.1), and you are upgrading a system with storage disks, you can skip to the [Map ports from node1 to node3](#).

Configure FC ports on node3

If node3 has FC ports, either onboard or on an add-on FC adapter, you must set port configurations on the node before you bring it into service because the ports are not preconfigured when the systems are shipped. If you don't configure the ports, 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 [Prepare the nodes for upgrade](#).

About this task

You can skip this section if your system does not have FC configurations. If your system has onboard UTA/UTA2 ports or a UTA/UTA2 card, you configure them in [Check and configure UTA/UTA2 ports on node3](#).



Enter the commands in this section at the cluster prompt.

Steps

- 1. Display information about all FC and converged network adapters on the system.

```
system node hardware unified-connect show
```

- 2. Compare the FC settings of node3 with the settings that you captured earlier from node1.
- 3. 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 9 .
Different from the ones that you captured on node1	Go to Step 4 .

- 4. Modify the FC ports on node3 as needed by entering one of the following commands:

- To program target ports:

```
system node hardware unified-connect modify -type \|-t target -adapter  
port_name
```

- To program initiator ports:

```
system node hardware unified-connect modify -type \|-t initiator -adapter  
port_name
```

-t is the FC4 type: target or initiator.

- 5. Verify the new settings by entering the following command and examining the output:

```
system node hardware unified-connect show
```

- 6. Exit Maintenance mode:

```
halt
```

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

8. Boot node3 at the boot environment prompt:

```
boot_ontap
```

9. Take one of the following actions:

- If node3 has a UTA/UTA2 card or UTA/UTA2 onboard ports, go to [Check and configure UTA/UTA2 ports on node3](#).
- If node3 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports, skip [Check and configure UTA/UTA2 ports on node3](#) and go to [Map ports from node1 to node3](#).

Check and configure UTA/UTA2 ports on node3

If node3 has onboard UTA/UTA2 ports or a UTA/UTA2 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 UTA/UTA2 ports.

About this task

If you want to use a Unified Target Adapter (UTA/UTA2) port for FC, you must first verify how the port is configured.



NetApp marketing materials might use the term UTA2 to refer to CNA adapters and ports. However, the CLI uses the term CNA.

You can use the `ucadmin show` command to verify the current port configuration:

```
*> ucadmin show
```

Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
0e	fc	target	-	initiator	offline
0f	fc	target	-	initiator	offline
0g	fc	target	-	initiator	offline
0h	fc	target	-	initiator	offline
1a	fc	target	-	-	online
1b	fc	target	-	-	online

6 entries were displayed.

UTA/UTA2 ports can be configured into native FC mode or UTA/UTA2 mode. FC mode supports FC initiator and FC target; UTA/UTA2 mode enables concurrent NIC and FCoE traffic sharing the same 10GbE SFP+ interface and supports FC targets.

UTA/UTA2 ports might be found on an adapter or on the controller, and have the following configurations, but you should check the configuration of the UTA/UTA2 ports on the node3 and change it, if necessary:

- UTA/UTA2 cards ordered when the controller is ordered are configured before shipment to have the personality you request.

- UTA/UTA2 cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard UTA/UTA2 ports on new controllers are configured before shipment to have the personality you request.



Enter the commands in this section at the cluster prompt unless directed to enter Maintenance mode.

Steps

1. Check the current port configuration by entering the following command on node3:

```
system node hardware unified-connect show
```

The system displays output similar to the following example:

```
cluster1::> system node hardware unified-connect show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
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.

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 to determine whether the UTA/UTA2 ports have the personality you want.
4. Take one of the following actions:

If the UTA/UTA2 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 13 and go to Step 14 .

5. If the system has storage disks and is running clustered Data ONTAP 8.3, boot node3 and enter maintenance mode:

```
boot_ontap maint
```

6. Verify the settings:

```
ucadmin show
```

7. Take one of the following actions:

If you are configuring...	Then...
Ports on a UTA/UTA2 card	Go to Step 8 .
Onboard UTA/UTA2 ports	Skip Step 8 and go to Step 9 .

8. If the adapter is in initiator mode, and if the UTA/UTA2 port is online, take the UTA/UTA2 port offline:

```
storage disable adapter adapter_name
```

Adapters in target mode are automatically offline in maintenance mode.

9. 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` is the FC4 type, `target` or `initiator`.



You must use the FC initiator for tape drives and MetroCluster configurations. You must use the FC target for SAN clients.

10. Stop the system:

```
halt
```

The system stops at the boot environment prompt.

11. Enter the following command:

```
boot_ontap
```

12. Verify the settings:

```
system node hardware unified-connect show
```

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

```
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.

13. Place any target ports online by entering the following command, once for each port:

```
network fcp adapter modify -node node_name -adapter adapter_name -state up
```

14. Cable the port.

Map ports from node1 to node3

You must 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*. (Go to [References](#) to link to the *Hardware Universe*). You use the information later in this section and in [Map ports from node2 to node4](#).

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

About this task

Port settings might vary, depending on the model of the nodes.

Steps

1. Perform the following steps to verify if the setup is a two-node switchless cluster:
 - a. Set the privilege level to advanced:

```
set -privilege advanced
```

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

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

For 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:

```
set -privilege admin
```

2. Make the following changes:

- a. Modify ports that will be part of Cluster broadcast domain:

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

This example adds Cluster port e1b on "node1":

```
network port modify -node node1 -port e1b -ipspace Cluster -mtu 9000
```

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

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

When all cluster LIFs are migrated and cluster communication is established, the cluster should come into quorum.

- c. Modify the home port of the Cluster LIFs:

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

- d. Remove the old ports from the Cluster broadcast domain:

```
network port broadcast-domain remove-ports -ipspace Cluster -broadcast
-domain Cluster -ports node1:port
```

- e. Display the health state of node1 and node3:

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

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

If your ONTAP version is...	Then...
9.8 to 9.11.1	Verify that the cluster LIFs are listening on port 7700: ::> network connections listening show -vserver Cluster

If your ONTAP version is...	Then...
9.12.1 or later	Skip this step and go to Step 3 .

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

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

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

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

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

3. Modify the broadcast domain memberships of physical ports hosting data LIFs.

- a. List the reachability status of all ports:

```
network port reachability show
```

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

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

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

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

- c. To enable ONTAP to complete the repair, wait for about a minute after running the `reachability repair` command on the last port.
- d. List all broadcast domains on the cluster:

```
network port broadcast-domain show
```


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

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

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

Add member ports to interface group ports:

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

Remove member ports from interface group ports:

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

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

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

Create VLAN ports:

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



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

4. If there are no VLANs configured on the system, go to [Step 5](#). If there are VLANs configured, restore displaced VLANs that were previously configured on ports that no longer exist or were configured on ports that were moved to another broadcast domain.

- a. Display the displaced VLANs:

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

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

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

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

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

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

```
network port reachability show
```

5. Beginning with ONTAP 9.8, ONTAP will automatically modify the home ports of LIFs if the ports are moved between broadcast domains during the network port reachability repair procedure. If a LIF's home

port was moved to another node, or is unassigned, that LIF will be presented as a displaced LIF. Restore the home ports of displaced LIFs whose home ports either no longer exist or were relocated to another node.

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

```
displaced-interface show
```

- b. Restore the home port of each LIF:

```
displaced-interface restore -vserver Vserver_name -lif-name LIF_name
```

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

```
displaced-interface show
```

When all ports are correctly configured and added to the correct broadcast domains, the `network port reachability show` command should report the reachability status as "ok" for all connected ports, and the status as "no-reachability" for ports with no physical connectivity. If any ports are reporting a status other than these two, repair the reachability as outlined in [Step 3](#).

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

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

```
network interface show -vserver Vserver_name -status-admin down
```

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

```
network interface show -vserver Vserver_name -status-oper down
```

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

```
network interface modify -vserver Vserver_name -lif LIF_name -home-port  
home_port
```



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

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

```
network interface revert *
```

Move NAS data LIFs owned by node1 from node2 to node3 and verify SAN LIFs on node3

Before you relocate aggregates from node2 to node3, you must move the NAS data LIFs belonging to node1 that are currently on node2 from node2 to node3. You also must 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 NAS 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. If the cluster is configured for SAN LIFs, record the SAN LIFs adapter and switch-port configuration information in this [worksheet](#) for use later in the procedure.
 - a. List the SAN LIFs on node2 and examine the output:

```
network interface show -data-protocol fc*
```

The system returns output similar to the following example:

```
cluster1::> net int show -data-protocol fc*
(network interface show)
Current Is      Logical      Status      Network      Current
Vserver        Interface    Admin/Oper  Address/Mask  Node
Port           Home
-----
svm2_cluster1
1b              true        lif_svm2_cluster1_340
                        up/up      20:02:00:50:56:b0:39:99
                        cluster1-01
1a              true        lif_svm2_cluster1_398
                        up/up      20:03:00:50:56:b0:39:99
                        cluster1-02
1a              true        lif_svm2_cluster1_691
                        up/up      20:01:00:50:56:b0:39:99
                        cluster1-01
1a              true        lif_svm2_cluster1_925
                        up/up      20:04:00:50:56:b0:39:99
                        cluster1-02
1b              true
4 entries were displayed.
```

- b. List the existing configurations and examine the output:

```
fcv adapter show -fields switch-port,fc-wwpn
```

The system returns output similar to the following example:

```

cluster1::> fcp adapter show -fields switch-port,fc-wwpn
(network fcp adapter show)
node          adapter  fc-wwpn                      switch-port
-----
cluster1-01   0a          50:0a:09:82:9c:13:38:00      ACME Switch:0
cluster1-01   0b          50:0a:09:82:9c:13:38:01      ACME Switch:1
cluster1-01   0c          50:0a:09:82:9c:13:38:02      ACME Switch:2
cluster1-01   0d          50:0a:09:82:9c:13:38:03      ACME Switch:3
cluster1-01   0e          50:0a:09:82:9c:13:38:04      ACME Switch:4
cluster1-01   0f          50:0a:09:82:9c:13:38:05      ACME Switch:5
cluster1-01   1a          50:0a:09:82:9c:13:38:06      ACME Switch:6
cluster1-01   1b          50:0a:09:82:9c:13:38:07      ACME Switch:7
cluster1-02   0a          50:0a:09:82:9c:6c:36:00      ACME Switch:0
cluster1-02   0b          50:0a:09:82:9c:6c:36:01      ACME Switch:1
cluster1-02   0c          50:0a:09:82:9c:6c:36:02      ACME Switch:2
cluster1-02   0d          50:0a:09:82:9c:6c:36:03      ACME Switch:3
cluster1-02   0e          50:0a:09:82:9c:6c:36:04      ACME Switch:4
cluster1-02   0f          50:0a:09:82:9c:6c:36:05      ACME Switch:5
cluster1-02   1a          50:0a:09:82:9c:6c:36:06      ACME Switch:6
cluster1-02   1b          50:0a:09:82:9c:6c:36:07      ACME Switch:7
16 entries were displayed

```

3. Take one of the following actions:

If node1...	Then...
Had interface groups or VLANs configured	Go to Step 4 .
Did not have interface groups or VLANs configured	Skip Step 4 and go to Step 5 .

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

- 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

```

- 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

```

- 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
```

5. Take one of the following actions:

If the cluster is configured for...	Then...
NAS	Complete Step 6 and Step 7 , skip Step 8, and complete Step 9 through Step 12 .
SAN	Disable all the SAN LIFs on the node to take them down for the upgrade: <pre>network interface modify -vserver vservice_name -lif LIF_name -home-node node_to_upgrade -home-port netport ifgrp -status-admin down</pre>

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

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

The following example adds port "e0a" on node "8200-1" 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 8200-1:e0a, 8060-1:e0i
```

7. Migrate each NAS 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
```

8. Make sure that the data migration is persistent:

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

9. Confirm that the SAN LIFs are on the correct ports on node3:

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

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

The system returns output similar to the following example:

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

Current	Is	Logical	Status	Network	Current
Vserver		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. Verify that the new and adapter and switch-port configurations are correct by comparing the output from the `fc adapter show` command with the configuration information that you recorded in the worksheet in [Step 2](#).

List the new SAN LIF configurations on node3:

```
fc adapter show -fields switch-port,fc-wwpn
```

The system returns output similar to the following example:

```
cluster1::> fcp adapter show -fields switch-port,fc-wwpn
(network fcp adapter show)
node          adapter fc-wwpn          switch-port
-----
cluster1-01 0a      50:0a:09:82:9c:13:38:00 ACME Switch:0
cluster1-01 0b      50:0a:09:82:9c:13:38:01 ACME Switch:1
cluster1-01 0c      50:0a:09:82:9c:13:38:02 ACME Switch:2
cluster1-01 0d      50:0a:09:82:9c:13:38:03 ACME Switch:3
cluster1-01 0e      50:0a:09:82:9c:13:38:04 ACME Switch:4
cluster1-01 0f      50:0a:09:82:9c:13:38:05 ACME Switch:5
cluster1-01 1a      50:0a:09:82:9c:13:38:06 ACME Switch:6
cluster1-01 1b      50:0a:09:82:9c:13:38:07 ACME Switch:7
cluster1-02 0a      50:0a:09:82:9c:6c:36:00 ACME Switch:0
cluster1-02 0b      50:0a:09:82:9c:6c:36:01 ACME Switch:1
cluster1-02 0c      50:0a:09:82:9c:6c:36:02 ACME Switch:2
cluster1-02 0d      50:0a:09:82:9c:6c:36:03 ACME Switch:3
cluster1-02 0e      50:0a:09:82:9c:6c:36:04 ACME Switch:4
cluster1-02 0f      50:0a:09:82:9c:6c:36:05 ACME Switch:5
cluster1-02 1a      50:0a:09:82:9c:6c:36:06 ACME Switch:6
cluster1-02 1b      50:0a:09:82:9c:6c:36:07 ACME Switch:7
16 entries were displayed
```



If a SAN LIF in the new configuration is not on an adapter that is still attached to the same switch-port, it might cause a system outage when you reboot the node.

- c. 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":

```
network interface modify -vserver vservice_name -lif LIF_name -status
-admin down
```

- ii. Remove the LIF from the port set:

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

- iii. Enter one of the following commands:

- Move a single LIF:

```
network interface modify -vserver vservice_name -lif LIF_name -home
-port new_home_port
```

- Move all the LIFs on a single nonexistent or incorrect port to a new port:

```
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
```



You must move SAN LIFs to a port that has the same link speed as the original port.

10. Modify the status of all LIFs to "up" so the LIFs can accept and send traffic on the node:

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

11. Enter the following command on either node 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
```

12. 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
```

13. Send a post-upgrade AutoSupport message to NetApp for node1:

```
system node autosupport invoke -node node3 -type all -message "node1  
successfully upgraded from platform_old to platform_new"
```

Worksheet: Information to record before moving NAS data LIFs to node3

To help verify that you have the correct configuration after moving SAN LIFs from node2 to node3, you can use the following worksheet to record the adapter and switch-port information for each LIF.

Record the LIF adapter information from the `network interface show -data-protocol fc*` command output and the switch-port information from the `fcport adapter show -fields switch-port,fc-wwpn` command output for node2.

After you complete the migration to node3, record the LIF adapter and switch-port information for the LIFs on node3 and verify that each LIF is still connected to the same switch-port.

Node2			Node3		
LIF	adapter	switch-port	LIF	adapter	switch-port

Node2			Node3		

Relocate non-root aggregates from node2 to node3

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



During this procedure, don't relocate aggregates from node3 to node2. Doing so results in aggregates being taken offline and a data outage for the aggregates that are relocated.

Steps

1. Verify that the partner system ID is set correctly on node3:

- a. Enter the advanced privilege level:

```
set -privilege advanced
```

- b. Show the partner system ID on node3:

```
ha interconnect config show -node <node3-node1>
```

The system displays output similar to the following example:

Show example

```
cluster::*> ha interconnect config show -node <node>
(system ha interconnect config show)

Node: node3-node1
Interconnect Type: RoCE
Local System ID: <node3-system-id>
Partner System ID: <node2-system-id>
Connection Initiator: local
Interface: external

Port    IP Address
----    -
e4a-17  0.0.0.0
e4b-18  0.0.0.0
```

2. If "Partner System ID" is incorrect for node3:

- a. Halt node3:

```
halt
```

- b. At the LOADER prompt, set the correct "partner-sysid" value.

The node3 "partner-sysid" is the system ID of node2, which you can find in the `ha interconnect config show` output in [Step 1](#).

- c. Save the settings:

```
saveenv
```

- d. At the LOADER prompt, boot node3 into the boot menu:

```
boot_ontap menu
```

- e. Log in to node3.

3. 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>"
```

4. 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.

5. Relocate the non-root aggregates:

- a. Set the privilege level to advanced:

```
set -privilege advanced
```

- b. List the aggregates that are owned by node2:

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

- c. Start aggregate relocation:

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



The command locates only non-root aggregates.

- d. 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 doesn't relocate any offline or restricted aggregates.

- e. Return to the admin privilege level:

```
set -privilege admin
```

6. Verify the relocation status of node2:

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


The output displays "Done" for an aggregate after it has been relocated.



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

7. Take one of the following actions:

If relocation of...	Then...
All aggregates was successful	Go to Step 8 .

If relocation of...	Then...
Any aggregates failed, or was vetoed	<p>a. Display a detailed status message:</p> <pre>storage aggregate show -instance</pre> <p>You can also check the EMS logs to see the corrective action that is needed.</p> <div style="display: flex; align-items: center; margin-top: 10px;">  <div> <p>The event log show command lists any errors that have occurred.</p> </div> </div> <p>b. Perform the corrective action.</p> <p>c. Set the privilege level to advanced:</p> <pre>set -privilege advanced</pre> <p>d. Relocate any failed or vetoed aggregates:</p> <pre>storage aggregate relocation start -node <node2> -destination <node3> -aggregate-list * -ndo-controllerupgrade true</pre> <p>e. When prompted, enter y.</p> <p>f. Return to the admin privilege level:</p> <pre>set -privilege admin</pre> <p>If necessary, you can force the relocation by using one of the following methods:</p> <ul style="list-style-type: none"> • By overriding veto checks: <pre>storage aggregate relocation start -override -vetoes true -ndo-controller-upgrade</pre> • By overriding destination checks: <pre>storage aggregate relocation start -override -destination-checks true -ndocontroller-upgrade</pre> <p>For more information about the storage aggregate relocation commands, go to References to link to <i>Disk and aggregate management with the CLI</i> and the <i>ONTAP 9 Commands: Manual Page Reference</i>.</p>

8. 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 <aggregate_name>
```

9. 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 <Vserver-name> -volume <volume-name>
```

10. Verify that node2 doesn't 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.

Move NAS data LIFs owned by node2 to node3

After you relocate the aggregates from node2 to node3, you need to move the NAS data LIFs owned by node2 to 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 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

1. List all the NAS 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
```

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...	Then...
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 .

- Take the following steps to migrate NAS data LIFs hosted on interface groups and VLANs on node2:
 - 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 Vserver_name -lif LIF_name -destination
-node node3 -destination-port netport|ifgrp
```

- 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 Vserver_name -lif LIF_name -home-node
node3 -homeport netport|ifgrp
```

- 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 Vserver_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 Vserver_name -lif LIF_name -home-node
node3 -homeport 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, add the ports to the broadcast domain:

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

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
```

6. Migrate each NAS 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. Verify that NAS 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 Vserver_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:

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

- b. Enter the following command and examine its output to determine 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: MAC_address
      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
```

Stage 4. Record information and retire node2

Record node2 information

Before you can shut down and retire node2, you must 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,nodemgmt,cluster-mgmt
```

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

```

Is	Logical	Status	Network	Current	Current
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

4 entries were displayed.



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 is displayed in the example output of [Step 1](#) in *Record node1 port information*.

2. Capture the information in the output to use in the section [Map 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.

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

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

(Mbps)					Speed	
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper
-----	-----	-----	-----	-----	-----	-----
node2						
	e0M	Default	IP_address	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
```

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 [Map ports from node2 to node4](#).

- Take one of the following actions:

If you...	Then...
Recorded NVRAM System ID number in Prepare the nodes for upgrade	Go to Retire node2 .
Did not record the NVRAM System ID number in Prepare the nodes for upgrade	Complete Step 9 and Step 10 and then go to the next section, Retire node2 .

- Display the attributes of node2:

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

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

- Record the NVRAM System ID to use in the section [Install and boot node4](#).

Retire node2

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

Steps

- Take one of the following actions:

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

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

```
set -privilege advanced
```

- 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
```

4. Check if node2 currently holds epsilon by entering the following command and examining its output:

```
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.



If you are upgrading an HA pair in a cluster with multiple HA pairs, you must move epsilon to the node of an HA pair that isn't undergoing a controller upgrade. For example, if you are upgrading nodeA/nodeB in a cluster with the HA pair configuration nodeA/nodeB and nodeC/nodeD, you must move epsilon to nodeC or nodeD.

5. If node2 holds epsilon, mark epsilon as false on the node so that it can be transferred to node3:

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

6. Transfer epsilon to node3 by marking epsilon true on node3:

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

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

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

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

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

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

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

```
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. Return to the admin level:

```
set -privilege admin
```

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

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

11. After node2 shuts down completely, remove it from the chassis or the rack. You can decommission node2 after the upgrade is completed. See [Decommission the old system](#).

Stage 5. Install and boot node4

Install and boot node4

You must 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 must netboot node4 if the ONTAP version on node4 is different to the ONTAP version on node2. After you install node4, boot it 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 by following the instructions in [Prepare for netboot](#)

However, you don't need to netboot node4 if the ONTAP version on node4 is the same or later than the ONTAP version on node2.



- For an AFF A800 or AFF C800 controller upgrade, you must ensure that all drives in the chassis are firmly seated against the midplane before removing node2. For more information, see [Replace the AFF A800 or AFF C800 controller modules](#).
- If you're upgrading a system with storage disks, you must complete this entire section and then proceed to the section [Set the FC or UTA/UTA2 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 Steps 2 and 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.

The following references help you make proper cable connections. Go to [References](#) to link to them.

- *Installation and Setup Instructions* for the node4 platform
- The appropriate disk shelf procedure
- The *HA pair management* documentation

Cable the following connections:

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



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.



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:
 - 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 finish.



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

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

```
halt
```

The system stops at the boot environment prompt.

9. Configure node4 for ONTAP:

```
set-defaults
```

10. If you have NetApp Storage Encryption (NSE) drives installed, perform the following steps:



If you have not already done so earlier in the procedure, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use.

- a. Set `bootarg.storageencryption.support` to `true` or `false`:

If the following drives are in use...	Then...
NSE drives that conform to FIPS 140-2 Level 2 self-encryption requirements	<code>setenv bootarg.storageencryption.support true</code>
NetApp non-FIPS SEDs	<code>setenv bootarg.storageencryption.support false</code>



You cannot mix FIPS drives with other types of drives on the same node or HA pair.

You can mix SEDs with non-encrypting drives on the same node or HA pair.

- b. Go to the special boot menu and select option (10) `Set Onboard Key Manager recovery secrets`.

Enter the passphrase and the backup information that you recorded earlier procedure. See [Manage authentication keys using the Onboard Key Manager](#).

11. If the version of ONTAP installed on node4 is the same or later than the version of ONTAP 9 installed on node2, enter the following command:

```
boot_ontap menu
```

12. Take one of the following actions:

If the system you are upgrading...	Then...
Does not have the correct or current ONTAP version on node4	Go to Step 13 .

If the system you are upgrading...	Then...
Has the correct or current version of ONTAP on node4	Go to Step 18 .

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



You must use the management port and IP address as the netboot connection. Do not use a data LIF IP address 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: <code>ifconfig e0M -auto</code>
Not running	<p>Manually configure the connection by entering the following command at the boot environment prompt:</p> <pre>ifconfig e0M -addr=filer_addr mask=netmask - gw=gateway dns=dns_addr domain=dns_domain</pre> <p><i>filer_addr</i> is the IP address of the storage system (mandatory). <i>netmask</i> is the network mask of the storage system (mandatory). <i>gateway</i> is the gateway for the storage system (mandatory). <i>dns_addr</i> is the IP address of a name server on your network (optional). <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> <div> <p>Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p> </div>

14. Perform netboot on node4:

For...	Then...
FAS/AFF8000 series systems	<pre>netboot http://<web_server_ip/path_to_webaccessible_directory> /netboot/kernel</pre>
All other systems	<pre>netboot http://<web_server_ip/path_to_webaccessible_directory> ontap_version>_image.tgz</pre>

The `<path_to_the_web-accessible_directory>` should lead to where you downloaded the `<ontap_version>_image.tgz` in [Step 1](#) in the section *Prepare for netboot*.



Do not interrupt the boot.

15. 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.



Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the incorrect image might install. This issue applies to all releases of ONTAP. The netboot procedure combined with option (7) Install new software wipes the boot media and places the same ONTAP version on both image partitions.

16. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL:

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

17. 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}
```

- 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.


18. Select maintenance mode 5 from the boot menu and enter `y` when you are prompted to continue with the boot.
19. Before continuing, go to [Set the FC or UTA/UTA2 configuration on node4](#) to make any necessary changes to the FC or UTA/UTA2 ports on the node. Make the changes recommended in those sections, reboot the node, and go into Maintenance mode.
20. Enter the following command and examine the output to find the system ID of node4:

```
disk show -a
```

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
-----
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)
.....
```

21. Reassign node2’s spares, disks belonging to the root, and any non-root aggregates that were not relocated to node3 earlier in section [Relocate non-root aggregates from node2 to node3](#):



If you have shared disks, hybrid aggregates, or both on your system, you must use the correct `disk reassign` command from the following table.

Disk type...	Run the command...
With shared disks	<pre>disk reassign -s node2_sysid -d node4_sysid -p node3_sysid</pre>
Without shared	<pre>disks disk reassign -s node2_sysid -d node4_sysid</pre>

For the `<node2_sysid>` value, use the information captured in [Step 10](#) of the *Record node2 information* section. For `node4_sysid`, use the information captured in [Step 23](#).



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
```

Enter `n` when asked to abort disk reassignment.

When you are asked to abort disk reassignment, you must answer a series of prompts as shown in the following steps:

- a. 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
```

- b. Enter `y` to continue.

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
```

- c. Enter `y` to allow disk ownership to be updated.

22. If you are upgrading from a system with external disks to a system that supports internal and external disks (A800 systems, for example), set `node4` as root to confirm that it boots from the root aggregate of `node2`.



Warning: You must perform the following substeps in the exact order shown; failure to do so might cause an outage or even data loss.

The following procedure sets `node4` to boot from the root aggregate of `node2`:

- a. Check the RAID, plex, and checksum information for the `node2` aggregate:

```
aggr status -r
```

- b. Check the overall status of the `node2` aggregate:

```
aggr status
```

- c. If necessary, bring the `node2` aggregate online:

```
aggr_online root_aggr_from_node2
```

d. Prevent the node4 from booting from its original root aggregate:

```
aggr offline root_aggr_on_node4
```

e. Set the node2 root aggregate as the new root aggregate for node4:

```
aggr options aggr_from_node2 root
```

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

```
ha-config show
```

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
```

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

If you have a MetroCluster configuration, use the following commands to correct the configuration:

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

```
ha-config modify chassis mcc.
```

24. Destroy the mailboxes on node4:

```
mailbox destroy local
```

25. Exit Maintenance mode:

```
halt
```

The system stops at the boot environment prompt.

26. On node3, check the system date, time, and time zone:

```
date
```

27. On node4, check the date at the boot environment prompt:

```
show date
```

28. If necessary, set the date on node4:

```
set date mm/dd/yyyy
```

29. On node4, check the time at the boot environment prompt:

```
show time
```

30. If necessary, set the time on node4:

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

31. Verify the partner system ID is set correctly as noted in [Step 19](#) under option.

```
printenv partner-sysid
```

32. If necessary, set the partner system ID on node4:

```
setenv partner-sysid node3_sysid
```

a. Save the settings:

```
saveenv
```

33. Enter the boot menu at the boot environment prompt:

```
boot_ontap menu
```

34. 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?:
```

35. Enter **y** at the prompt.

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



The system might reboot twice before displaying the mismatch warning.

36. Confirm the mismatch.

The node might complete one round of rebooting before booting normally.

37. Log in to node4.

Set the FC or UTA/UTA2 configuration on node4

If node4 has onboard FC ports, onboard unified target adapter (UTA/UTA2) ports, or a UTA/UTA2 card, you must configure the settings before completing the rest of the procedure.

About this task

You might need to complete [Configure FC ports on node4](#) or [Check and configure UTA/UTA2 ports on node4](#), or both sections.

If node4 does not have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card (for example, AFF and FAS systems introduced beginning with ONTAP 9.15.1), and you are upgrading a system with storage disks, you can skip to [Map ports from node2 to node4](#).

Configure FC ports on node4

If node4 has FC ports, either 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 [Prepare the nodes for upgrade](#).

About this task

You can skip this section if your system does not have FC configurations. If your system has onboard UTA/UTA2 ports or a UTA/UTA2 adapter, you configure them in [Check and configure UTA/UTA2 ports on node4](#).



Enter the commands in this section at the cluster prompt.

Steps

1. Display information about all FC and converged network adapters on the system:

```
system node hardware unified-connect show
```

2. Compare the FC settings on the new nodes with the settings that you captured earlier from the original node.
3. Modify the FC ports on node4 as needed:
 - To program target ports:

```
system node hardware unified-connect modify -type \|-t target -adapter  
port_name
```

- To program initiator ports:

```
system node unified-connect modify type \|-t initiator -adapter port_name  
  
-type is the FC4 type, target or initiator.
```

4. Verify the new settings by entering the following command and examining the output:

```
system node unified-connect show
```

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 the original nodes	Go to Step 9 .
Different from the ones that you captured on the original nodes	Go to Step 6 .

6. Exit Maintenance mode:

```
halt
```

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

8. Boot node4 by entering the following command at the boot environment prompt:

```
boot_ontap
```

9. Take one of the following actions:

- Go to [Check and configure UTA/UTA2 ports on node4](#) if node4 has a UTA/UTA2A card or UTA/UTA2 onboard ports.
- Skip the section and go to [Map ports from node2 to node4](#) if node4 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports.

Check and configure UTA/UTA2 ports on node4

If node4 has onboard UTA/UTA2 ports or a UTA/UTA2A 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 UTA/UTA2 ports.

About this task

UTA/UTA2 ports can be configured into native FC mode or UTA/UTA2A mode. FC mode supports FC initiator and FC target; UTA/UTA2 mode enables concurrent NIC and FCoE traffic to share the same 10GbE SFP+ interface and supports FC target.



NetApp marketing materials might use the term UTA2 to refer to CNA adapters and ports. However, the CLI uses the term CNA.

UTA/UTA2 ports might be on an adapter or on the controller with the following configurations:

- UTA/UTA2 cards ordered at the same time as the controller are configured before shipment to have the personality you requested.
- UTA/UTA2 cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard UTA/UTA2 ports on new controllers are configured (before shipment) to have the personality you requested.

However, you can check the configuration of the UTA/UTA2 ports on node4 and change it, if necessary.



Enter the commands in this section at the cluster prompt unless directed to enter Maintenance mode. If you have a MetroCluster FC system, you must be in Maintenance mode to configure UTA/UTA2 ports.

Steps

1. Check how the ports are currently configured on node4:

```
system node hardware unified-connect show
```

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 UTA/UTA2 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 to Step 13 .

5. If the system has storage disks and is running Data ONTAP 8.3, boot node4 and enter maintenance mode:

```
boot_ontap maint
```

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

```
ucadmin show
```

7. Take one of the following actions:

If you are configuring...	Then...
Ports on a UTA/UTA2A card	Go to Step 8 .
Onboard UTA/UTA2 ports	Skip Step 8 and go to Step 9 .

8. If the adapter is in initiator mode, and if the UTA/UTA2 port is online, take the UTA/UTA2 port offline:

```
storage disable adapter adapter_name
```

Adapters in target mode are automatically offline in Maintenance mode.

9. 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 UTA.

- -t is the FC4 type: target or initiator.



You must use FC initiator for tape drives and the FC target for SAN clients.

- If the system has storage disks, enter the following command:

```
halt
```

The system stops at the boot environment prompt.

- Enter the following command:

```
boot_ontap
```

- If the system has storage disks, enter the following command:

```
system node hardware unified-connect show
```

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.

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

```
network fcp adapter modify -node node_name -adapter adapter_name -state up
```

- Cable the port.

Map ports from node2 to node4

You must 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

You must already have information about the ports on the new nodes, to access this information refer to [References](#) to link to the *Hardware Universe*. You use the information later in this section.

The software configuration of node4 must match the physical connectivity of node4, 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.

Steps

1. Perform the following steps to verify if the setup is a two-node switchless cluster:

- a. Set the privilege level to advanced:

```
set -privilege advanced
```

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

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

For 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 using the following command:

```
set -privilege admin
```

2. Make the following changes:

- a. Modify ports that will be part of Cluster broadcast domain:

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

This example adds Cluster port "e1b" on "node2":

```
network port modify -node node2 -port e1b -ipspace Cluster -mtu 9000
```

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

```
network interface migrate -vserver vserver_name -lif lif_name source-node
node2 -destination-node node2 -destination-port port_name
```

When all cluster LIFs are migrated and cluster communication is established, the cluster should come into quorum.

- c. Modify the home port of the Cluster LIFs:

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

- d. Remove the old ports from the Cluster broadcast domain:

```
network port broadcast-domain remove-ports -ipspace Cluster -broadcast
-domain Cluster -ports node2:port
```

- e. Display the health state of node2/node4:

```
cluster show -node node2 -fields health
```

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

If your ONTAP version is...	Then...
9.8 to 9.11.1	Verify that the cluster LIFs are listening on port 7700: ::> network connections listening show -vserver Cluster
9.12.1 or later	Skip this step and go to Step 3 .

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

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

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

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

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

3. Modify the broadcast domain memberships of physical ports hosting data LIFs.

- a. List the reachability status of all ports:

```
network port reachability show
```

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

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

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

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

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

- d. List all broadcast domains on the cluster:

```
network port broadcast-domain show
```

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

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

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

Add member ports to interface group ports:

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

Remove member ports from interface group ports:

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

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

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

Create VLAN ports:

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



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

- 4. If there are no VLANs configured on the system, go to [Step 5](#). If there are VLANs configured, restore displaced VLANs that were previously configured on ports that no longer exist or were configured on ports that were moved to another broadcast domain.

- a. Display the displaced VLANs:

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

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

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

`destination_port`

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

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

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

```
network port reachability show
```

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

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

```
displaced-interface show
```

- b. Restore the home port of each LIF:

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

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

```
displaced-interface show
```

When all ports are correctly configured and added to the correct broadcast domains, the `network port reachability show` command should report the reachability status as `ok` for all connected ports, and the status as `no-reachability` for ports with no physical connectivity. If any ports are reporting a status other than these two, repair the reachability as outlined in [Step 3](#).

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

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

```
network interface show -vserver vservice_name -status-admin down
```

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

```
network interface show -vserver vservice_name -status-oper down
```

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

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



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

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

```
network interface revert *
```

Move NAS data LIFs owned by node2 from node3 to node4 and verify SAN LIFs on node4

After mapping ports from node2 to node4 and before you relocate node2 aggregates from node3 to node4, you must move the NAS data LIFs owned by node2 currently on node3 from node3 to node4. You also must 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 NAS 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. If the cluster is configured for SAN LIFs, record the SAN LIFs and existing configuration information in this [worksheet](#) for use later in the procedure.
 - a. List the SAN LIFs on node3 and examine the output:

```
network interface show -data-protocol fc*
```

The system returns output similar to the following example:

```

cluster1::> net int show -data-protocol fc*
(network interface show)

```

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			
-----	-----	-----	-----	-----
-----	-----			
svm2_cluster1	lif_svm2_cluster1_340	up/up	20:02:00:50:56:b0:39:99	cluster1-01
1b	true			
	lif_svm2_cluster1_398	up/up	20:03:00:50:56:b0:39:99	cluster1-02
1a	true			
	lif_svm2_cluster1_691	up/up	20:01:00:50:56:b0:39:99	cluster1-01
1a	true			
	lif_svm2_cluster1_925	up/up	20:04:00:50:56:b0:39:99	cluster1-02
1b	true			

4 entries were displayed.

b. List the existing configurations and examine the output:

```
fcip adapter show -fields switch-port,fc-wwpn
```

The system returns output similar to the following example:

```
cluster1::> fcp adapter show -fields switch-port,fc-wwpn
(network fcp adapter show)
node          adapter  fc-wwpn                      switch-port
-----
cluster1-01   0a         50:0a:09:82:9c:13:38:00     ACME Switch:0
cluster1-01   0b         50:0a:09:82:9c:13:38:01     ACME Switch:1
cluster1-01   0c         50:0a:09:82:9c:13:38:02     ACME Switch:2
cluster1-01   0d         50:0a:09:82:9c:13:38:03     ACME Switch:3
cluster1-01   0e         50:0a:09:82:9c:13:38:04     ACME Switch:4
cluster1-01   0f         50:0a:09:82:9c:13:38:05     ACME Switch:5
cluster1-01   1a         50:0a:09:82:9c:13:38:06     ACME Switch:6
cluster1-01   1b         50:0a:09:82:9c:13:38:07     ACME Switch:7
cluster1-02   0a         50:0a:09:82:9c:6c:36:00     ACME Switch:0
cluster1-02   0b         50:0a:09:82:9c:6c:36:01     ACME Switch:1
cluster1-02   0c         50:0a:09:82:9c:6c:36:02     ACME Switch:2
cluster1-02   0d         50:0a:09:82:9c:6c:36:03     ACME Switch:3
cluster1-02   0e         50:0a:09:82:9c:6c:36:04     ACME Switch:4
cluster1-02   0f         50:0a:09:82:9c:6c:36:05     ACME Switch:5
cluster1-02   1a         50:0a:09:82:9c:6c:36:06     ACME Switch:6
cluster1-02   1b         50:0a:09:82:9c:6c:36:07     ACME Switch:7
16 entries were displayed
```

3. Take one of the following actions:

If node2...	Description
Had interface groups or VLANs configured	Go to Step 4 .
Did not have interface groups or VLANs configured	Skip Step 4 and go to Step 5 .

4. Take the following steps to migrate any NAS data LIFs hosted on interface groups and VLANs that originally were on node2 from node3 to node4.

- 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 vservice_name -lif lif_name -destination
-node node4 -destination-port netport|ifgrp
```

- 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 vservice_name -lif datalif_name -home-node
node4 home-port netport|ifgrp
```

- 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 vservice_name -lif datalif_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 vservice_name -lif datalif_name -home-node
node4 home-port netport|ifgrp
```

5. Take one of the following actions:

If the cluster is configured for...	Then...
NAS	Complete Step 6 through Step 9 , skip Step 10 , and complete Step 11 through Step 14 .
SAN	Skip Step 6 through Step 9 , and complete Step 10 through Step 14 .
Both NAS and SAN	Complete Step 6 through Step 14 .

6. 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
```

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
```

7. Migrate each NAS 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
```

8. Make sure that the data migration is persistent:

```
network interface modify -vserver vservice_name -lif datalif_name -home-port
netport|ifgrp
```

9. 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
```

The following example shows the output of the `network port show` command with some LIFs up and others down:

```
cluster::> network port show
```

(Mbps)						Speed
Node	Port	IPspace	Broadcast Domain	Link	MTU	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.						

10. 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:

- a. Enter the advanced privilege level by entering the following command:

```
set -privilege advanced
```

- b. Enter the following command, once for each old network port:

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

- c. Return to the admin level by entering the following command:

```
set -privilege admin
```

11. Confirm that the SAN LIFs are on the correct ports on node4 by completing the following substeps:

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

```
network interface show -data-protocol iscsi|fc -home-node node4
```

The system returns output similar to the following example:

```
cluster::> network interface show -data-protocol iscsi|fc -home-node
node4
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			
-----	-----	-----	-----	-----
-----	----			
vs0				
	a0a	up/down	10.63.0.53/24	node4
a0a	true			
	data1	up/up	10.63.0.50/18	node4
e0c	true			
	rads1	up/up	10.63.0.51/18	node4
e1a	true			
	rads2	up/down	10.63.0.52/24	node4
e1b	true			
vs1				
	lif1	up/up	172.17.176.120/24	node4
e0c	true			
	lif2	up/up	172.17.176.121/24	node4

- b. Verify that the new adapter and switch-port configurations are correct by comparing the output from the `fc -adapter show` command with the new configuration information that you recorded in the worksheet in [Step 2](#).

List the new SAN LIF configurations on node4:

```
fc -adapter show -fields switch-port,fc-wwpn
```

The system returns output similar to the following example:

```
cluster1::> fcp adapter show -fields switch-port,fc-wwpn
(network fcp adapter show)
node          adapter  fc-wwpn          switch-port
-----
cluster1-01   0a          50:0a:09:82:9c:13:38:00  ACME Switch:0
cluster1-01   0b          50:0a:09:82:9c:13:38:01  ACME Switch:1
cluster1-01   0c          50:0a:09:82:9c:13:38:02  ACME Switch:2
cluster1-01   0d          50:0a:09:82:9c:13:38:03  ACME Switch:3
cluster1-01   0e          50:0a:09:82:9c:13:38:04  ACME Switch:4
cluster1-01   0f          50:0a:09:82:9c:13:38:05  ACME Switch:5
cluster1-01   1a          50:0a:09:82:9c:13:38:06  ACME Switch:6
cluster1-01   1b          50:0a:09:82:9c:13:38:07  ACME Switch:7
cluster1-02   0a          50:0a:09:82:9c:6c:36:00  ACME Switch:0
cluster1-02   0b          50:0a:09:82:9c:6c:36:01  ACME Switch:1
cluster1-02   0c          50:0a:09:82:9c:6c:36:02  ACME Switch:2
cluster1-02   0d          50:0a:09:82:9c:6c:36:03  ACME Switch:3
cluster1-02   0e          50:0a:09:82:9c:6c:36:04  ACME Switch:4
cluster1-02   0f          50:0a:09:82:9c:6c:36:05  ACME Switch:5
cluster1-02   1a          50:0a:09:82:9c:6c:36:06  ACME Switch:6
cluster1-02   1b          50:0a:09:82:9c:6c:36:07  ACME Switch:7
16 entries were displayed
```



If a SAN LIF in the new configuration is not on an adapter that is still attached to the same switch-port, it might cause a system outage when you reboot the node.

- c. 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:

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

- ii. Remove the LIF from the port set:

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

- iii. Enter one of the following commands:

- Move a single LIF:

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

- Move all the LIFs on a single nonexistent or incorrect port to a new port:

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



You must move SAN LIFs to a port that has the same link speed as the original port.

12. 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 vserver_name -home-port port_name -home-node
node4 lif lif_name -status-admin up
```

13. 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
```

14. 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
```

Worksheet: Information to record before moving NAS data LIFs to node4

To help verify that you have the correct configuration after moving SAN LIFs from node3 to node4, you can use the following worksheet to record the adapter and switch-port information for each LIF.

Record the LIF adapter information from the `network interface show -data-protocol fc*` command output and the switch-port information from the `fcport adapter show -fields switch-port, fc-wwpn` command output for node3.

After you complete the migration to node4, record the LIF adapter and switch-port information for the LIFs on node4 and verify that each LIF is still connected to the same switch-port.

Node3			Node4		
LIF	adapter	switch-port	LIF	adapter	switch-port

Relocate node2 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
```

- b. 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](#).

- c. 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.

- d. Return to the admin level:

```
set -privilege admin
```

- 3. Check the relocation status:

```
storage aggregate relocation show -node node3
```

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



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 .

If relocation of...	Then...
Any aggregates failed, or were 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:</p> <pre>set -privilege advanced</pre> <p>d. Relocate any failed or vetoed aggregates:</p> <pre>storage aggregate relocation start -node node3 destination node4 -aggregate-list aggr_name1, aggr_name2... ndo-controller-upgrade true</pre> <p>The aggregate list is the list of failed or vetoed aggregates.</p> <p>e. When prompted, enter y.</p> <p>f. Return to the admin level by entering the following command:</p> <pre>set -privilege admin</pre> <p>If necessary, you can force the relocation using one of the following methods:</p> <ul style="list-style-type: none"> • Overriding veto checks: <pre>storage aggregate relocation start -override -vetoes -ndo-controller-upgrade</pre> • Overriding destination checks: <pre>storage aggregate relocation start -override -destination-checks -ndocontroller-upgrade</pre> <p>For more information about storage aggregate relocation commands refer to References to link to <i>Disk and aggregate management with the CLI</i> and the <i>ONTAP 9 Commands: Manual Page Reference</i>.</p>

5. Verify that all node2 non-root aggregates are online and their state on node4:

```
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:

```
volume show -node node4 -state offline
```

8. If any volumes are offline on node4, bring them online:

```
volume online -vserver vserver-name -volume volume_name
```

9. Send a post-upgrade AutoSupport message to NetApp for node4:

```
system node autosupport invoke -node node4 -type all -message "node2  
successfully upgraded from platform_old to platform_new"
```

Stage 6. Complete the upgrade

Manage authentication using KMIP servers

With ONTAP 9.5 and later, you can use Key Management Interoperability Protocol (KMIP) servers to manage authentication keys.

Steps

1. Add a new controller:

```
security key-manager setup -node new_controller_name
```

2. Add the key manager:

```
security key-manager -add key_management_server_ip_address
```

3. Verify that the key management servers are configured and available to all nodes in the cluster:

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

4. Restore the authentication keys from all linked key management servers to the new node:

```
security key-manager restore -node new_controller_name
```

Confirm that the new controllers are set up correctly

To confirm correct setup, you enable the HA pair. You 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 confirm 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:

```
storage failover show
```


The following example shows the output of the command when storage failover is enabled:

```
cluster::> storage failover show
```

Node	Partner	Takeover Possible	State Description
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: <code>cluster ha modify -configured true</code>
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:

```
storage aggregate show -owner-name <node3>
storage aggregate show -owner-name <node4>
```

8. Determine whether any volumes are offline:

```
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 [Step 19 \(d\)](#) in *Prepare the nodes for upgrade*, and bring online any of the offline volumes, as required, by entering the following command, once for each volume:

```
volume online -vserver <vserver_name> -volume <volume_name>
```

10. Install new licenses for the new nodes by entering the following command 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 multiple licenses at once, each license key separated by a comma.

11. If self-encrypting drives are being used in the configuration and you have set the `kmip.init.maxwait` variable to `off` (for example, in [Step 16](#) of *Install and boot node3*), you must unset the variable:

```
set diag; systemshell -node node_name -command sudo kenv -u -p  
kmip.init.maxwait
```

12. 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.

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

```
system license show
```

You can compare the output with the output that you captured in [Step 29](#) of *Prepare the nodes for upgrade*.

14. Configure the SPs by performing the following command on both nodes:

```
system service-processor network modify -node <node_name>
```

Go to [References](#) to link to the *System Administration Reference* for information about the SPs and the *ONTAP 9 Commands: Manual Page Reference* for detailed information about the `system service-`

processor network modify command.

15. If you want to set up a switchless cluster on the new nodes, go to [References](#) to link to the *Network Support Site* and follow the instructions in *Transitioning to a two-node switchless cluster*.

After you finish

If Storage Encryption is enabled on node3 and node4, complete the steps in [Set up Storage Encryption on the new controller module](#). Otherwise, complete the steps in [Decommission the old system](#).

Set 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. You must 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 the previous step 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 the 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
```

Set up NetApp Volume or Aggregate Encryption on the new controller module

If the replaced controller or the high-availability (HA) partner of the new controller uses NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE), you must configure the new controller module for NVE or NAE.

About this task

This procedure includes steps that are performed on the new controller module. You must 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 key query -node node
```

2. Add the key management servers listed in the previous step to the key management server list in the new controller:

- a. Add the key management server by using the following command:

```
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 by using the following command:

```
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 by using the following command:

```
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:

For...	Use this command...
External key manager	<pre>security key-manager external restore</pre> <p>This command needs the OKM passphrase</p>
Onboard Key Manager (OKM)	<pre>security key-manager onboard sync</pre>

For more information, see the Knowledge Base article [How to restore external key manager server configuration from the ONTAP boot menu](#).

After you finish

Check if any volumes were taken offline because authentication keys were not available or External Key Management servers could not be reached. Bring those volumes back online by using the `volume online` command.

Decommission 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. Refer to [References](#) to link to the *NetApp Support Site* and log in.
2. Select **Products > My Products** from the menu.
3. On the **View Installed Systems** page, choose which **Selection Criteria** you want to use to display information about your system.

You can choose one of the following to locate your system:

- Serial Number (located on the back of the unit)
- Serial Numbers for My Location

4. Select **Go!**

A table displays cluster information, including the serial numbers.

5. Locate the cluster in the table and select **Decommission this system** from the Product Tool Set drop-down menu.

Resume SnapMirror operations

You can resume SnapMirror transfers that were quiesced before upgrade and resume the SnapMirror relationships. The updates are on schedule after the upgrade is completed.

Steps

1. Verify the SnapMirror status on the destination:

```
snapmirror show
```

2. Resume the SnapMirror relationship:

```
snapmirror resume -destination-vserver vservers_name
```

Troubleshoot

Aggregate relocation failures

Aggregate relocation (ARL) might fail at different points during the upgrade.

Check for aggregate relocation failure

During the procedure, ARL might fail in Stage 2, Stage 3, or Stage 5.

Steps

1. Enter the following command and examine the output:

```
storage aggregate relocation show
```

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, refer to [References](#) to link to the *ONTAP 9 Commands: Manual Page Reference*.

Aggregates originally on node1 are owned by node4 after completion of the upgrade

At the end of the upgrade procedure, node3 must 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 correctly, 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 [Prepare the nodes for 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 [Prepare the nodes for upgrade](#) and note any aggregates that were not correctly relocated.
3. Relocate the aggregates left behind on node4:

```
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 looks 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 *Relocate non-root aggregates from node1 to node2* again.
2. Continue with the node-pair upgrade procedure.

Node1 crashes after Stage 2 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 fails 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 gives back all of node2's aggregates.

Steps

1. Complete [Step 1](#) in the section *Relocate 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 the rest of the node pair upgrade procedure.

Reboots, panics, or power cycles during Stage 3

Failures 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 5](#) in the section *Relocate non-root 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 does 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 5](#) again in the section *Relocate non-root 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 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 *Relocate 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 *Relocate 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 during 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 Possible	State Description
node1	node2	false	Unknown
node2	node1	false	Node owns partner aggregates as part of the non-disruptive head upgrade procedure. Takeover is not possible: Storage failover is disabled.

The following example shows the output of the `storage failover show` command after a reboot or panic:

```
cluster::> storage failover show
```

Node	Partner	Takeover Possible	State Description
node1	node2	-	Unknown
node2	node1	false	Waiting for node1, Partial giveback, Takeover is not possible: Storage failover is disabled

Although the output says that a node is in partial giveback and that storage failover is disabled, you can disregard this message.

Steps

No action is required; continue with the node-pair upgrade procedure.

LIF migration failure

After you migrate LIFs, they might not come online after migration in Stage 2, Stage 3, or

Stage 5.

Steps

1. Verify that the 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.

2. Check the physical connectivity of the network cable if the physical state of the port is "down".

LIFs are on invalid ports after upgrade

After the upgrade is completed, the FC logical interfaces (LIFs) might be left on incorrect ports if you have a MetroCluster configuration. You can perform a resync operation to reassign the LIFs to the correct ports.

Step

1. Enter the `metrocluster vserver resync` command to reallocate the LIFs to the correct ports.

```
metrocluster vserver resync -vserver vserver_name fcp-mc.headupgrade.test.vs
```

References

When performing the procedures in this content, you might need to consult reference content or go to reference websites.

Reference content

Content specific to this upgrade are listed in the table below.

Content	Description
Administration overview with the CLI	Describes how to administer ONTAP systems, shows you how to use the CLI interface, how to access the cluster, how to manage nodes, and much more.
Decide whether to use System Manager or the ONTAP CLI for cluster setup	Describes how to set up and configure ONTAP.
Disk and aggregate management with the CLI	Describes how to manage ONTAP physical storage using the CLI. It shows you how to create, expand, and manage aggregates, how to work with Flash Pool aggregates, how to manage disks, and how to manage RAID policies.
HA pair management	Describes how to install and manage high-availability clustered configurations, including storage failover and takeover/giveback.
Logical storage management with the CLI	Describes how to efficiently manage your logical storage resources, using volumes, FlexClone volumes, files, and LUNs, FlexCache volumes, deduplication, compression, qtrees, and quotas.

Content	Description
MetroCluster Management and Disaster Recovery	Describes how to perform MetroCluster switchover and switchback operations, both in planned maintenance operations, or in the event of a disaster.
MetroCluster Upgrade and Expansion	Provides procedures for upgrading controller and storage models in the MetroCluster configuration, transitioning from a MetroCluster FC to a MetroCluster IP configuration, and expanding the MetroCluster configuration by adding additional nodes.
Network Management	Describes how to configure and manage physical and virtual network ports (VLANs and interface groups), LIFs, routing, and host-resolution services in clusters; optimize network traffic by load balancing; and monitor the cluster by using SNMP.
ONTAP 9.0 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.0 commands.
ONTAP 9.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.1 commands.
ONTAP 9.2 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.2 commands.
ONTAP 9.3 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.3 commands.
ONTAP 9.4 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.4 commands.
ONTAP 9.5 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.5 commands.
ONTAP 9.6 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.6 commands.
ONTAP 9.7 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.7 commands.
ONTAP 9.8 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.8 commands.
ONTAP 9.9.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.9.1 commands.
ONTAP 9.10.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.10.1 commands.
SAN management with the CLI	Describes how to configure and manage LUNs, igroups, and targets using the iSCSI and FC protocols, and namespaces and subsystems using the NVMe/FC protocol.
SAN configuration reference	Contains information about FC and iSCSI topologies and wiring schemes.
Upgrade by moving volumes or storage	Describes how to quickly upgrade controller hardware in a cluster by moving storage or volumes. Also describes how to convert a supported model to a disk shelf.

Content	Description
Upgrade ONTAP	Contains instructions for downloading and upgrading ONTAP.
Use "system controller replace" commands to upgrade controller hardware introduced in ONTAP 9.15.1 and later	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers introduced in ONTAP 9.15.1 and later by using "system controller replace" commands.
Use "system controller replace" commands to upgrade controller models in the same chassis	Describes the aggregate relocation procedures needed to non-disruptively upgrade a system, keeping the old system chassis and disks.
Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers running ONTAP 9.8 by using "system controller replace" commands.
Use aggregate relocation to manually upgrade controller hardware running ONTAP 9.8 or later	Describes the aggregate relocation procedures needed to perform manual non-disruptive controller upgrades running ONTAP 9.8 or later.
Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.5 to ONTAP 9.7	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers running ONTAP 9.5 to ONTAP 9.7 by using "system controller replace" commands.
Use aggregate relocation to manually upgrade controller hardware running ONTAP 9.7 or earlier	Describes the aggregate relocation procedures needed to perform manual non-disruptive controller upgrades running ONTAP 9.7 or earlier.

Reference sites

The [NetApp Support Site](#) also contains documentation about network interface cards (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.

Access [ONTAP 9 documentation](#).

Access the [Active IQ Config Advisor](#) tool.

Manually upgrade controller hardware running ONTAP 9.7 or earlier

Learn about this ARL upgrade procedure

This procedure describes how to upgrade the controller hardware using manual aggregate relocation (ARL) on systems running ONTAP 9.7 or earlier.

You can use this ARL procedure if you are performing one of the following upgrades:

- FAS system to FAS system
- AFF system to AFF system

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 confirm 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.



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*.

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.

Important information:

- This procedure is complex and assumes that you have advanced ONTAP administration skills. You also must read and understand the [guidelines for upgrading controllers with ARL](#) and the [ARL upgrade workflow](#) before beginning the upgrade.
- 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.
- 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 high-availability (HA) pair in the cluster.
- This procedure applies to MetroCluster four-node and eight-node configurations running ONTAP 9.5 and earlier. For MetroCluster configurations running ONTAP 9.6 and later, go to [References](#) to link to *Using "system controller replace" Commands to Upgrade Controller Hardware Running ONTAP 9.5 to ONTAP 9.7*.

Decide whether to use this aggregate relocation procedure

This procedure describes how to upgrade controller hardware using manual aggregate relocation (ARL) on systems running ONTAP 9.7 or earlier. You should only use this complex procedure if you're an experienced ONTAP administrator.

Use this content under the following circumstances:

- You don't 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.
- Your controllers are running ONTAP 9.7 or earlier.
- You have a system that uses Fabric MetroCluster 4-node and 8-node configurations running ONTAP 9.5 or earlier.



- If you're upgrading a system by swapping controller modules in the same chassis, such as AFF A800 or AFF C800, NetApp strongly recommends using the upgrade procedure that [upgrades controller models using ARL, keeping the existing system chassis and disks](#). This ARL procedure includes the steps that ensure the internal disks remain secure in the chassis when you remove and install the controllers during the upgrade procedure.

[Learn about the supported system upgrade combinations using ARL, keeping the existing system chassis and disks.](#)

- You can use NetApp Storage Encryption (NSE), NetApp Volume Encryption (NVE), and NetApp Aggregate Encryption (NAE) with this procedure.

If you prefer a different method of upgrading the controller hardware and are willing to do volume moves, refer to [References](#) to link to *Upgrade by moving volumes or storage*.

Refer to [References](#) to link to the *ONTAP 9 Documentation Center* where you can access ONTAP 9 product documentation.

Choose a different hardware upgrade procedure

- [Review the alternative ARL methods available for upgrading controller hardware.](#)
- If you prefer a different method of upgrading the controller hardware and are willing to do volume moves, refer to [References](#) to link to *Upgrade by moving volumes or storage*.

Related information

Refer to [References](#) to link to the *ONTAP 9 Documentation*.

ARL upgrade workflow

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.

Upgrade the node pair


To upgrade the node pair, you must 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.

ARL upgrade sequence overview

During the procedure, you upgrade the original controller hardware with the replacement controller hardware, one controller at a time, taking advantage of the HA pair 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:

Stage	Description
Stage 1: Prepare for upgrade	<p>During Stage 1, if required, you confirm that internal disk drives do not contain root aggregates or data aggregates, prepare the nodes for the upgrade, and run a series of prechecks. If required, you rekey disks for Storage Encryption and prepare to netboot the new controllers.</p> <p>Aggregate ownership at the end of Stage 1:</p> <ul style="list-style-type: none"> Node1 is the home owner and current owner of the node1 aggregates. Node2 is the home owner and current owner of the node2 aggregates.

Stage	Description
Stage 2: Retire node1	<p>During Stage 2, you relocate non-root aggregates from node1 to node2 and move non-SAN data LIFs owned by node1 to node2, including failed or vetoed aggregates. You record the necessary node1 information for use later in the procedure and then retire node1.</p> <p>Aggregate ownership at the end of Stage 2:</p> <ul style="list-style-type: none"> • Node1 is the home owner of node1 aggregates. • Node2 is the current owner of node1 aggregates. • Node2 is the home owner and current owner of node2 aggregates.
Stage 3: Install and boot node3	<p>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.</p> <p>Aggregate ownership at the end of Stage 3:</p> <ul style="list-style-type: none"> • Node2 is the home owner of node2 aggregates but not the current owner. • Node3 is the home owner and current owner of aggregates originally belonging to node1. • Node2 is the home owner and current owner of aggregates belonging to node2 but not the home owner.
Stage 4: Retire node2	<p>During Stage 4, you record the necessary node2 information for use later in the procedure and then retire node2.</p> <p>No changes occur in aggregate ownership.</p>
Stage 5: Install and boot node4	<p>During Stage 5, you install and boot node4, map the cluster and node-management ports from node2 to node4, verify the node4 installation, and move data LIFs and SAN LIFs belonging to node2 from node3 to node4. You also relocate node2 aggregates from node3 to node4 and move the data node2 NAS LIFs from node3 to node4.</p> <p>Aggregate ownership at the end of Stage 5:</p> <ul style="list-style-type: none"> • Node3 is the home owner and current owner of the aggregates that originally belonged to node1. • Node4 is the home owner and current owner of aggregates that originally belonged to node2.

Stage	Description
Stage 6: Complete the upgrade	<p>During Stage 6, you confirm that the new nodes are set up correctly and set up Storage Encryption or NetApp Volume Encryption if the new nodes are encryption-enabled. You should also decommission the old nodes and resume SnapMirror operations.</p> <div>  <p>The storage virtual machine (SVM) disaster recovery updates will not be interrupted as per the schedules assigned.</p> </div> <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 9.0 to 9.7 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.
 - An AFF system to an AFF system.



Before performing an AFF system upgrade, you must upgrade ONTAP to release versions 9.3P12, 9.4P6 or 9.5P1 or later. These release levels are required for a successful upgrade.

- For some ARL controller upgrades you can use temporary cluster ports on the replacement controller for the upgrade. For example, if you upgrade from an AFF A300 to an AFF A400 system, depending on the AFF A400 configuration, you can use any of the two mezzanine ports or add a four-port 10GbE network interface card to provide temporary cluster ports. After you complete a controller upgrade using temporary cluster ports, you can nondisruptively migrate clusters to 100GbE ports on the replacement controller.
- If you are using ONTAP 9.6P11, 9.7P8, or later releases, it is recommended to enable Connectivity, Liveliness, and Availability Monitor (CLAM) takeover to return the cluster into quorum when certain node failures occur. The `kernel-service` command requires advanced privilege level access. For more information, see: [NetApp KB Article SU436: CLAM takeover default configuration changed](#).
- Controller upgrade using ARL is supported on systems configured with SnapLock Enterprise and SnapLock Compliance volumes.

You must verify whether the ARL can be performed on the original and replacement controllers. You must 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. To access this information, refer to [References](#) to link to the *Hardware Universe*. 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 must 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. For more information about cluster mixing rules, refer to [References](#) to link to the *Hardware Universe*.



Both systems are either high-availability (HA) or non-HA. 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, contact technical support.



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. Refer to [References](#) to link to 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 a version of ONTAP from ONTAP 9.0 to ONTAP 9.7.

For information on controller upgrades for systems running Data ONTAP operating in 7-Mode, refer to [References](#) to link to the *NetApp Support Site*.

- To replacement controllers that do not support the disk shelves connected to the original controllers.

For disk-support information, refer to [References](#) to link to the *Hardware Universe*.

- 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, refer to [References](#) to link to *Upgrade by moving volumes or storage* and go to the procedure *Upgrading a pair of nodes running clustered Data ONTAP by moving volumes*.



If you want to upgrade ONTAP on nodes in a cluster, refer to [References](#) to link to *Upgrade ONTAP*.

Assumptions and terminology

This document is written with the following assumptions:

- 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 because this procedure assumes that the replacement controller hardware is new and has not been used. You must contact technical support if the replacement controller hardware was previously used, especially if the controllers were running Data ONTAP in 7-Mode.

- You read and understand the guidelines for upgrading the pair of nodes.



Do not try to clear the NVRAM contents. If you 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 partner LIFs for each storage virtual machine (SVM), on the HA pair. If you do not have local and partner 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 9.7 or earlier

Some features require 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 and will continue to work. However, using unlicensed features on the controller might put you out of compliance with your license agreement, so you must install the new license key or keys for the new controller after the upgrade is complete.

All license keys are 28 uppercase alphabetic characters in length. Refer to [References](#) to link to the *NetApp Support Site* where you can obtain new 28-character license keys for ONTAP 9.7. or earlier. The keys are available in the *My Support* section under *Software licenses*. If the site does not have the license keys you need, contact your NetApp sales representative.

For detailed information about licensing, go to [References](#) to link to the *System Administration Reference*.

Storage Encryption

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

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.

Troubleshoot

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.

If any problems occur while upgrading the controllers, refer to the [Troubleshoot](#) section. The information about failures that can occur is listed by the phase of the procedure in the [ARL upgrade sequence](#).

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

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 must 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

Go to the [References](#) section to access the list of reference documents required for this upgrade.

Worksheet: Information to collect before and during controller upgrade

You should gather certain information to support upgrading the original nodes. This information includes node IDs, port and LIF details, licensing keys, and IP addresses.

You can use the following worksheet to record the 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>Prepare the nodes for the upgrade</i>	Stage 3: <i>Install and boot node3</i> Stage 5: <i>Install and boot node4</i> Stage 6: <i>Decommission 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>Prepare the nodes for the upgrade</i>	Throughout the procedure to verify that aggregates and volumes remain online except during brief relocation	

Information needed	When collected	When used	Collected Information
Output of commands network port vlan show and network port ifgrp show	Stage 1: <i>Prepare the nodes for the upgrade</i>	Stage 3: <i>Map ports from node1 to node3</i> Stage 5: <i>Map ports from node2 to node4</i>	
(SAN environments only) Default configuration of FC ports	Stage 1: <i>Prepare the nodes for the upgrade</i>	When configuring FC ports on the new nodes	
IP address of SPs	Stage 1: <i>Prepare the nodes for the upgrade</i>	Stage 6: <i>Confirm that the new controllers are set up correctly</i>	
License keys	Stage 1: <i>Prepare the nodes for the upgrade</i>	Stage 6: <i>Confirm that the new controllers are set up correctly</i>	
IP address for the External Key Management server	Stage 1: <i>Rekey disks for Storage Encryption</i>	Stage 6: <i>Set 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>Prepare to netboot</i>	Stage 3: <i>Install and boot node3</i> Stage 5: <i>Install and boot node4</i>	
Non-SAN data LIFs owned by node1	Stage 2: <i>Move nonSAN data LIFs owned by node1 to node2</i>	Later in the section	
Cluster, intercluster, node-management, cluster-management, and physical ports	Stage 2: <i>Record node1 information</i>	Stage 3: <i>Install and boot node3</i> Stage 3: <i>Map ports from node1 to node3</i>	
Ports on new nodes	Stage 3: <i>Map ports from node1 to node3</i>	Later in the section and in the section <i>Map ports from node2 to node4</i>	
Available ports and broadcast domains on node3	Stage 3: <i>Map 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>Move nonSAN data LIFs owned by node2 to node3</i>	Later in the section	

Information needed	When collected	When used	Collected Information
Cluster, intercluster, node-management, cluster-management, and physical ports	Stage 4: <i>Record node2 information</i>	Stage 5: <i>Install and booting node4</i> Stage 5: <i>_ Map ports from node2 to node4_</i>	
Cluster network ports on node4	Stage 5: <i>Map ports from node2 to node4</i>	Later in the section	
Available ports and broadcast domains on node4	Stage 5: <i>Map 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>Set up Storage Encryption on the new nodes</i>	Later in the section	

Reconfigure the FC switch layout for ONTAP 9.1 or later

Reconfigure the FC switch layout for ONTAP 9.1 or later

If your existing FC switch layout was configured prior to ONTAP 9.1, you must reconfigure the port layout and apply the latest Reference Configuration Files (RCFs). This procedure applies only to MetroCluster FC configurations.

Before you begin

You must identify the FC switches present in the fabric domain.

You need the admin password and access to an FTP or SCP server.

About this task

You must perform this task if your existing FC switch layout was configured prior to ONTAP 9.1 and you are upgrading to a platform model supported in ONTAP 9.1 or later. It is *not* required if you are upgrading from an existing switch layout that was configured for ONTAP 9.1 or later.

This procedure is nondisruptive and takes approximately four hours to complete (excluding rack and stack) when disks are zeroed.

Steps

1. [Send a custom AutoSupport message prior to reconfiguring switches](#)
2. [Verify the health of the MetroCluster configuration](#)
3. [Check for MetroCluster configuration errors](#)
4. [Persistently disable the switches](#)
5. [Determine the new cabling layout](#)
6. [Apply RCF files and recable the switches](#)
7. [Persistently enable the switches](#)

8. [Verify switchover, healing, and switchback](#)

Send a custom AutoSupport message prior to reconfiguring switches

Before reconfiguring your switches, you must issue an AutoSupport message to notify NetApp technical support that maintenance is underway. Informing technical support that maintenance is underway prevents them from opening a case on the assumption that a disruption has occurred.

About this task

This task must be performed on each MetroCluster site.

Steps

1. Log in to the cluster.
2. Invoke an AutoSupport message indicating the start of the maintenance:

```
system node autosupport invoke -node * -type all -message MAINT=maintenance-  
window-in-hours
```

The `maintenance-window-in-hours` value specifies the length of the maintenance window, with a maximum of 72 hours. If the maintenance is completed before the time has elapsed, you can invoke an AutoSupport message indicating the end of the maintenance period:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

3. Repeat these steps on the partner site.

Verify the health of the MetroCluster configuration

You should check the health of the MetroCluster configuration to verify correct operation.

Steps

1. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::> metrocluster check run
```

```
Last Checked On: 10/1/2017 16:03:37
```

Component	Result
nodes	ok
lifs	ok
config-replication	ok
aggregates	ok

4 entries were displayed.

Command completed. Use the "metrocluster check show -instance" command or sub-commands in "metrocluster check" directory for detailed results. To check if the nodes are ready to do a switchover or switchback operation, run "metrocluster switchover -simulate" or "metrocluster switchback -simulate", respectively.

2. Verify that there are no health alerts:

```
system health alert show
```

Check for MetroCluster configuration errors

You can use the Active IQ Config Advisor tool available from the NetApp Support Site to check for common configuration errors.

If you do not have a MetroCluster configuration, you can skip this section.

About this task

Active IQ Config Advisor is a configuration validation and health check tool. You can deploy it at both secure sites and non-secure sites for data collection and system analysis.



Support for Config Advisor is limited, and available only online.

1. Download the [Active IQ Config Advisor](#) tool.
2. Run Active IQ Config Advisor, reviewing the output and following its recommendations to address any issues.

Persistently disable the switches

You must disable the switches in the fabric persistently so that you can modify its configuration.

About this task

You disable the switches by running the commands on the switch command line; the commands used for this are not ONTAP commands.

Step

Persistently disable the switch:

- For Brocade switches, use the `switchCfgPersistentDisable` command.
- For Cisco switches, use the `suspend` command.

The following command disables a Brocade switch persistently:

```
FC_switch_A_1:admin> switchCfgPersistentDisable
```

The following command disables a Cisco switch:

```
vsan [vsna #] suspend
```

Determine the new cabling layout

You must determine the cabling for the new controller modules and any new disk shelves to the existing FC switches.

About this task

This task must be performed at each MetroCluster site.

Step

Use the *Fabric-attached MetroCluster Installation and Configuration* content to determine the cabling layout for your switch type, using the port usage for an eight-node MetroCluster configuration. The FC switch port usage must match the usage described in the content so that the Reference Configuration Files (RCFs) can be used.

Go to [References](#) to link to the *Fabric-attached MetroCluster Installation and Configuration* content.



If your environment cannot be cabled in a way that RCFs can be used, contact technical support. Do not use this procedure if the cabling cannot use RCFs.

Apply RCF files and recable the switches

You must apply the appropriate reference configuration files (RCFs) to reconfigure your switches to accommodate the new nodes. After you apply the RCFs, you can recable the switches.

Before you begin

The FC switch port usage must match the usage described in the *Fabric-attached MetroCluster Installation and Configuration* content so that the RCFs can be used. Go to [References](#) to link to the *Fabric-attached MetroCluster Installation and Configuration* content.

Steps

1. Go to the [MetroCluster RCF downloads](#) page and select the RCFs for your switch configuration.

You must use the RCFs that match your switch models.

2. Install the FC switch RCFs by selecting the procedure that matches your switch models and following the installation instructions:
 - [Install a Brocade FC switch RCF](#)
 - [Install a Cisco FC switch RCF](#)
3. Verify that the switch configuration is saved.
4. Cable both of the FC-to-SAS bridges to the FC switches, using the cabling layout you created in [Determine the new cabling layout](#).
5. Verify that the ports are online:
 - For Brocade switches, use the `switchshow` command.
 - For Cisco switches, use the `show interface brief` command.
6. Cable the FC-VI ports from the controllers to the switches.
7. From the existing nodes, verify that the FC-VI ports are online:

```
metrocluster interconnect adapter show
```

```
metrocluster interconnect mirror show
```

Persistently enable the switches

You must enable the switches in the fabric persistently.

Step

Persistently enable the switch:

- For Brocade switches, use the `switchCfgPersistentenable` command.

```
FC_switch_A_1:admin> switchCfgPersistentenable
```

- For Cisco switches, use the `no suspend` command.

```
vsan [vsna #]no suspend
```

Verify switchover, healing, and switchback

You should verify the switchover, healing, and switchback operations of the MetroCluster configuration.

Step

Refer to [References](#) to link to the *MetroCluster Management and Disaster Recovery* content and follow the procedures for negotiated switchover, healing, and switchback.

Stage 1. Prepare for upgrade

Determine 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 confirm 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 [Prepare the nodes for 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
```

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)
...

```

3. 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.

4. Enter the following command on both controllers:

```
aggr status -r
```

The system displays the aggregate status of the node, as shown in the partial output in 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
...

```



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

5. 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 the previous step, "aggr2" uses internal drives, as indicated by the shelf ID of "0".

6. Enter the following command on both controllers:

```
aggr status -v
```

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,
        64-bit                raidsize=14, 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
        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.

- 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.
- Exit the nodeshell by entering the following command on each controller:

```
exit
```

- Take one of the following actions:

If the controllers....	Then...
Do not contain any aggregates on internal disk drives	Continue with this procedure.

If the controllers....	Then...
Contain aggregates but no volumes on the internal disk drives	<p>Continue with this procedure.</p> <div>  <p>Before you continue, you must place the aggregates offline, and then destroy the aggregates on the internal disk drives. Refer to References to link to the <i>Disk and aggregate management with the CLI</i> content for information about managing aggregates.</p> </div>
Contain non-root volumes on the internal drives	<p>Continue with this procedure.</p> <div>  <p>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. Refer to References to link to the <i>Disk and aggregate management with the CLI</i> content for information about moving volumes.</p> </div>
Contain root volumes on the internal drives	<p>Do not continue with this procedure.</p> <p>You can upgrade the controllers by referring to References to link to the <i>NetApp Support Site</i> and using the procedure <i>Upgrading the controller hardware on a pair of nodes running clustered Data ONTAP by moving volumes</i>.</p>
Contain non-root volumes on the internal drives and you cannot move the volumes to external storage	<p>Do not continue with this procedure.</p> <p>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>. Refer to References to link to the <i>NetApp Support Site</i> where you can access this procedure.</p>

Prepare the nodes for upgrade

Before you can replace the original nodes, you must confirm 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 must collect information about the original nodes and, if the cluster is in a SAN environment, confirm that all the nodes in the cluster are in quorum.

Steps

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

Refer to [References](#) to link to *HA pair management* and follow the *Best practices for HA pairs* section. 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.



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

For information about unlocking the user accounts, refer to [References](#) to link to the *System Administration Reference*.

- b. Refer to [References](#) to link to the *NetApp Support Site* and download the Performance and Statistics Collector (Perfstat Converged).

The 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. Refer to [References](#) to link to the *NetApp Support Site* and open a support case on the NetApp Support Site.

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 must physically check node3 and node4 to see if the NVMEM or NVRAM batteries are charged. For information about the LEDs for the model of node3 and node4, refer to [References](#) to link to the *Hardware Universe*.



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 must netboot the new controllers after you install them. For instructions on how to upgrade ONTAP, refer to [References](#) to link 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 node2:

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

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

```
cluster::> network interface show -role cluster
```

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

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. Confirm 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
```

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. For detailed information about the commands, refer to [References](#) to link to *HA pair management*.

9. Confirm that neither node1 nor node2 owns the 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
```

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,homename,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 .

If the command in Step 9...	Then...
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
```

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. Confirm 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
```

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, refer to [References](#) to link to *Disk and aggregate management with the CLI*, *Logical storage management with the CLI*, and *HA pair management* to configure storage for the HA pair.

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

```
cluster show
```

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. Set the privilege level to advanced:

```
set -privilege advanced
```

15. Confirm that node1 and node2 are running the same ONTAP release:

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

The following example shows the output of the command:

```
cluster::*> system node image show -node node1,node2 -iscurrent true
```

Node	Image	Is Default	Is Current	Version	Install 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.

16. 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
```

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
```

Current Is	Logical	Status	Network	Current	
Vserver	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.

17. 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 *
```

For detailed information about the `network interface revert` command, refer to [References](#) to link to the *ONTAP 9 Commands: Manual Page Reference*.

18. 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 *Disk and aggregate management with the CLI*. (Refer to [References](#) to link to *Disk and aggregate management with the CLI*.)

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



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
```



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
```



You can use the information to identify parts or accessories that you might want to transfer to node3 or node4.

- 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
```



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
```



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.

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

```
network port ifgrp show
```

```
network port vlan 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.

21. Complete the following substeps on both node1 and node2 to confirm 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.

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

```
cluster::> network interface failover-groups show
```

Vserver	Group	Targets
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.

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

```
cluster::> network port vlan show
```

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.
22. Take one of the following actions:

If interface groups or VLANs are...	Then...
On node1 or node2	Complete Step 23 and Step 24 .
Not on node1 or node2	Go to Step 24 .

23. 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 vservice_name -data-protocol iscsi|fc
```

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 Vserver8970 -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 `network interface show` command with the `-data-protocol nfs|cifs` parameters.

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 vs1 -data-protocol iscsi|fc
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
vs1	vs1_lif1	up/down	172.17.176.20/24	node1	0d	true

24. 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
```

The system displays a message similar to the following example:

```
cluster::*> cluster kernel-service show
```

Master Node	Cluster Node	Quorum Status	Availability Status	Operational 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, refer to [References](#) to link to the *System Administration Reference*.

d. Return to the administrative privilege level:

```
set -privilege admin
```

25. Take one of the following actions:

If the cluster...	Then...
Has SAN configured	Go to Step 26 .
Does not have SAN configured	Go to Step 29 .

26. 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|fc -home-node node_name
```

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 iscsi|fc
          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 following command:

```
network interface show -instance -data-protocol iscsi|fc
```

27. 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
```

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.

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. Confirm that the SP status is `online`.
c. Confirm 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, refer to [References](#) to link to 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 *
```

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 at the *NetApp Support Site*. Refer to [References](#) to link to *NetApp Support Site*.

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
```

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...	Go to Step 34 .

If the original system...	Then...
Does not have AutoSupport enabled...	<p>Enable AutoSupport by following the instructions in the <i>System Administration Reference</i>. (Refer to References to link to the <i>System Administration Reference</i>.)</p> <p>Note: 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.</p>

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, refer to [References](#) to link to 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"
```



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.

36. If your system uses self-encrypting drives, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use on the HA pair that you are upgrading. ONTAP software supports two types of self-encrypting drives:

- FIPS-certified NetApp Storage Encryption (NSE) SAS or NVMe drives
- Non-FIPS self-encrypting NVMe drives (SED)



You cannot mix FIPS drives with other types of drives on the same node or HA pair.

You can mix SEDs with non-encrypting drives on the same node or HA pair.

[Learn more about supported self-encrypting drives.](#)

Manage authentication keys using the Onboard Key Manager

You can use the Onboard Key Manager (OKM) to manage authentication keys. If you have the OKM set up, you must record the passphrase and backup material before beginning the upgrade.

Steps

1. Record the cluster-wide passphrase.

This is the passphrase that was entered when the OKM was configured or updated using the CLI or REST API.

2. Back up the key-manager information by running the `security key-manager onboard show-backup` command.

Quiesce the SnapMirror relationships

Before you netboot the system, you must confirm that all the SnapMirror relationships are quiesced. When a SnapMirror relationship is quiesced, it remains quiesced across reboots and failovers.

Steps

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

```
snapmirror show
```



If the status is `Transferring`, you must abort those transfers:

```
snapmirror abort -destination-vserver vservers name
```

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

2. Quiesce all relationships between the cluster:

```
snapmirror quiesce -destination-vserver *
```

Prepare for netboot

After you physically rack node3 and node4 later in the procedure, you might need to netboot them. The term *netboot* means you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you must put a copy of the ONTAP 9 boot image onto a web server that the system can access.

Before you begin

- Verify that you can access a HTTP server with the system.
- Refer to [References](#) to link to the *NetApp Support Site* and download the necessary system files for your platform and the correct version of ONTAP.

About this task


You must netboot the new controllers if they do not have the same 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 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 [Stage 3: Install and boot node3](#).

Steps

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

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

You will use information in the directories in [Stage 3](#).

Stage 2. Relocate and retire node1

Relocate non-root aggregates from node1 to node2

Before you can replace node1 with node3, you must 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. Set the privilege level to advanced:

```
set -privilege advanced
```

b. Enter the following command:

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

c. 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.

d. 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.



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:

If relocation...	Then..
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. Relocate any failed or vetoed aggregates: <pre>storage aggregate relocation start -node <i>node1</i> - destination <i>node2</i> -aggregate-list * -ndo -controller-upgrade true</pre> </p> <p>d. When prompted, enter <i>y</i>.</p> <p>e. Return to the admin level: <pre>set -privilege admin</pre> If necessary, you can force the relocation using one of the following methods:</p> <ul style="list-style-type: none"> ◦ Override veto checks: <pre>storage aggregate relocation start -override -vetoes true -ndo-controller-upgrade</pre> ◦ Override destination checks: <pre>storage aggregate relocation start -override -destination-checks true -ndo-controller -upgrade</pre> <p>Refer to References to link to the <i>Disk and aggregate management with the CLI</i> content 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:

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

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,

normal
aggr_2      825.0GB 825.0GB      0% online      1 node2
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
```

5. 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 vserver-name -volume volume-name
```

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

6. 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.
```

7. Verify that node1 does not own any non-root aggregates that are online:

```
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.

Move NAS data LIFs owned by node1 to node2

Before you can replace node1 with node3, you must move the NAS 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

1. List all the NAS data LIFs hosted on node1 by entering the following command and capturing the output:

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

The system displays the NAS data LIFs on node1, as shown in the following example:

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

Is	Logical	Status	Network	Current	Current
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					
	rads2	up/down	10.63.0.52/24	node1	e1b
true					
vs1					
	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...	Then...
Has interface groups of 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` command to display information about the network ports attached to VLANs, and use the `network port ifgrp show` command to display information about the port interface groups.

3. Take the following steps to migrate any NAS 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 Vserver_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 Vserver_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 Vserver-name -lif LIF_name -home-node node_to_upgrade -home-port netport ifgrp -status -admin down</pre>

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

```
network interface migrate -vserver Vserver-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 Vserver-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 Vserver-name -lif LIF_name  
-destination-node node_name -destination-port port_name
```

```
network interface modify -vserver Vserver-name -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 Vserver-name -lif LIF_name -home-node  
nodename -status-admin up
```



For MetroCluster configurations, 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 storage virtual machine (SVM). Enter the following command from the corresponding source SVM on the remote site to reallocate the destination LIF to an appropriate port:

```
metrocluster vsync resync -vserver Vserver_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_name -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
```

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
```

Record node1 information

Before you can shut down and retire node1, you must 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
```

The system displays output similar to the following example:

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

Vserver	Destination	Gateway	Metric
-----	-----	-----	-----
iscsi vsserver	0.0.0.0/0	10.10.50.1	20
node1	0.0.0.0/0	10.10.20.1	10
....			
node2	0.0.0.0/0	192.169.1.1	20

2. Enter the following command and capture its output:

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

The system displays output similar to the following example:

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

Vserver	State	Domains	Name Servers
node 1 2 10.10.60.10,	enabled	alpha.beta.gamma.netapp.com	
10.10.60.20 vs_base1 10.10.60.10,	enabled	alpha.beta.gamma.netapp.com, beta.gamma.netapp.com,	
10.10.60.20 ...			
...			
vs_peer1 10.10.60.10,	enabled	alpha.beta.gamma.netapp.com, gamma.netapp.com	
10.10.60.20			

- 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
```

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
```

Current Is	Logical	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	
vserver1	cluster mgmt	up/up	192.168.x.xxx/24	node1	e0c
true					
node1	intercluster	up/up	192.168.x.xxx/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.x.xxx/24	node1	e0c
true					

5 entries were displayed.



Your system might not have intercluster LIFs.

- Capture the information in the output of the command in [Step 3](#) to use in the section [Map ports from node1 to node3](#).

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

- Enter the following command on node1:

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

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

```
sti8080mcc-htp-008::> network port show -node sti8080mcc-htp-008 -type physical
```

Node: sti8080mcc-htp-008

Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Oper	Health Status	Ignore Health Status
e0M	Default	Mgmt	up	1500	auto/1000	healthy	false
e0a	Default	Default	up	9000	auto/10000	healthy	false
e0b	Default	-	up	9000	auto/10000	healthy	false
e0c	Default	-	down	9000	auto/-	-	false
e0d	Default	-	down	9000	auto/-	-	false
e0e	Cluster	Cluster	up	9000	auto/10000	healthy	false
e0f	Default	-	up	9000	auto/10000	healthy	false
e0g	Cluster	Cluster	up	9000	auto/10000	healthy	false
e0h	Default	Default	up	9000	auto/10000	healthy	false

9 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
```

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 [Map ports from node1 to node3](#).

10. Take one of the following actions:

If you...	Then...
Recorded the NVRAM System ID number in the section Prepare the nodes for the upgrade .	Go on to the next section, Retire node1 .
Did not record the NVRAM System ID number in the section Prepare the nodes for the upgrade	Complete Step 11 and Step 12 and then continue to Retire node1 .

11. Enter the following command on either controller:

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

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 [Install and boot node3](#).

Retire node1

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

Steps

1. Verify the number of nodes in the cluster:

```
cluster show
```

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



If you do not disable storage failover, a controller upgrade failure can occur which can disrupt data access and lead to data loss.

3. Verify that storage failover was disabled:

```
storage failover show
```

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 [Troubleshoot](#) section.

5. Take one of the following actions:

If the cluster is...	Then...
A two-node cluster	Go to Step 6 .
A cluster with more than two nodes	Go to Step 8 .

6. Access the advanced privilege level on either node:

```
set -privilege advanced
```

7. 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](#).

8. 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. Refer to [References](#) to link to the *System Administration Reference* for more information.



If you have a four-node cluster, epsilon might be on a node in a different HA pair in the cluster.

If you are upgrading an HA pair in a cluster with multiple HA pairs, you must move epsilon to the node of an HA pair that isn't undergoing a controller upgrade. For example, if you are upgrading nodeA/nodeB in a cluster with the HA pair configuration nodeA/nodeB and nodeC/nodeD, you must move epsilon to nodeC or nodeD.

The following example shows that node1 holds epsilon:

```
cluster::*> cluster show
```

Node	Health	Eligibility	Epsilon
node1	true	true	true
node2	true	true	false

9. 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
```

10. Transfer epsilon to node2 by marking epsilon true on node2:

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

11. Verify that the change to node2 occurred:

```
cluster show
```

```
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.

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

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



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

Enable Switchless Cluster: false/true
```

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

13. Return to the admin level:

```
set -privilege admin
```

14. Halt node1 from the node1 prompt:

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



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.

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

The node stops at the boot environment prompt.

16. 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 [Decommission the old system](#).

Stage 3. Install and boot node3

Install and boot node3

You must install node3 in the rack, transfer node1's connections to node3, boot node3, and install ONTAP. You must also 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 must netboot node3 if it doesn't have the same version of ONTAP 9 that is installed on node1. After you install node3, boot it 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. See [Prepare for netboot](#).

However, you don't need to netboot node3 if it has the same or a later version of ONTAP 9 that is installed on node1.



- For an AFF A800 or AFF C800 controller upgrade, you must ensure that all drives in the chassis are firmly seated against the midplane before removing node1. For more information, see [Replace the AFF A800 or AFF C800 controller modules](#).
- If you are upgrading a system with storage disks, you must complete this entire section and then go to [Configure FC ports on node3](#) and [Check and configure UTA/UTA2 ports on node3](#), 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 must put node3 into its own rack space, preferably close to the location of node1.

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



If you're upgrading to a system with both nodes in the same chassis, install node4 and node3 in the chassis. If you don't install both nodes in the same chassis, when you boot node3, it behaves as if it were in a dual-chassis configuration, and when you boot node4, the interconnect between the nodes doesn't come up.

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

The following references help you make proper cable connections. Go to [References](#) to link to them.

- *Installation and Setup Instructions* for the node3 platform
- The appropriate disk shelf procedure
- The *HA pair management* documentation

Cable the following connections:

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



You might 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 must move the FC-VI cable connections from node1 to node3. If the new host does not have an FC-VI card, you might need to move the FC-VI card.

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 later.



When you boot node3, you might see the following warning 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:

- 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 complete.




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

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

```
halt
```

The system stops at the boot environment prompt.

7. 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 8 .
Single-chassis configuration (with controllers in the same chassis)	<ol style="list-style-type: none">Switch the console cable from node3 to node4.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.  Leave node4 at the boot environment prompt; you will return to node4 in Install and boot node4.If you see the warning message displayed in Step 4, follow the instructions in Step 5Switch the console cable back from node4 to node3.Go to Step 8.

8. Configure node3 for ONTAP:

```
set-defaults
```

9. If you have NetApp Storage Encryption (NSE) drives installed, perform the following steps:



If you have not already done so earlier in the procedure, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use.

- a. Set `bootarg.storageencryption.support` to `true` or `false`:

If the following drives are in use...	Then...
NSE drives that conform to FIPS 140-2 Level 2 self-encryption requirements	<code>setenv bootarg.storageencryption.support true</code>
NetApp non-FIPS SEDs	<code>setenv bootarg.storageencryption.support false</code>



You cannot mix FIPS drives with other types of drives on the same node or HA pair.

You can mix SEDs with non-encrypting drives on the same node or HA pair.

- b. Contact NetApp Support for assistance with restoring the onboard key management information.

10. If the version of ONTAP installed on node3 is the same or later than the version of ONTAP 9 installed on node1, list and reassign disks to the new node3:

```
boot_ontap
```



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 might result in service outages or data loss. Contact technical support if the replacement controller was previously used, especially if the controllers were running ONTAP running in 7-Mode.

11. Press CTRL-C to display the boot menu.


12. Take one of the following actions:

If the system you are upgrading...	Then...
Does <i>not</i> have the correct or current ONTAP version on node3	Go to Step 13 .
Has the correct or current version of ONTAP on node3	Go to Step 18 .

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



You must use the management port and IP as the netboot connection. Do not use a data LIF IP or else 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: <code>ifconfig e0M -auto</code>
Not running	<p>Manually configure the connection by entering the following command at the boot environment prompt: <code>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></code></p> <p><i>filer_addr</i> is the IP address of the storage system (mandatory). <i>netmask</i> is the network mask of the storage system (mandatory). <i>gateway</i> is the gateway for the storage system (mandatory). <i>dns_addr</i> is the IP address of a name server on your network (optional). <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> <div>  <p>Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p> </div>

14. Perform netboot on node3:

For...	Then...
FAS/AFF8000 series systems	<code>netboot</code> <code>http://<web_server_ip>/<path_to_webaccessible_directory>/netboot/kernel</code>
All other systems	<code>netboot</code> <code>http://<web_server_ip>/<path_to_webaccessible_directory>/<ontap_version>_image.tgz</code>

The `<path_to_the_web-accessible_directory>` leads to where you downloaded the `<ontap_version>_image.tgz` in [Step 1](#) in the section *Prepare for netboot*.



Do not interrupt the boot.

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

This menu option downloads and installs the new 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 ONTAP, and not upgrades of controllers.



Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the wrong image might install. This issue applies to all releases of ONTAP. The netboot procedure combined with option (7) *Install new software* wipes the boot media and places the same ONTAP version ONTAP on both image partitions.

16. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the following URL:

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

17. 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}
```

- 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.

18. Select **(5) Maintenance mode boot** by entering `5`, and then enter `y` when prompted to continue with the boot.
19. Before continuing, go to [Set the FC or UTA/UTA2 configuration on node3](#) to make any necessary changes to the FC or UTA/UTA2 ports on the node.

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

20. Find the system ID of node3:

```
disk show -a
```

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  HOME      DR
HOME                                NUMBER
-----
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)
.....
```



You might 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.

21. Reassign node1's spare disks, any disks belonging to the root, and any non-root aggregates that were not relocated to node2 earlier in [Relocate non-root aggregates from node1 to node2](#).

Enter the appropriate form of the `disk reassign` command based on whether your system has shared disks:



If you have shared disks, hybrid aggregates, or both on your system, you must use the correct `disk reassign` command from the following table.

If disk type is...	Then run the command...
With shared disks	<code>disk reassign -s <i>node1_sysid</i> -d <i>node3_sysid</i> -p <i>node2_sysid</i></code>
Without shared disks	<code>disk reassign -s <i>node1_sysid</i> -d <i>node3_sysid</i></code>

For the `node1_sysid` value, use the information captured in [Record node1 information](#). To obtain the value for `node3_sysid`, use the `sysconfig` command.



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)?
```

22. 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)?
```

23. 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)?
```

24. Enter *y*.

25. If you are upgrading from a system with external disks to a system that supports internal and external disks (AFF A800 systems, for example), set the node1 aggregate as root to confirm that node3 boots from the root aggregate of node1.



Warning: You must perform the following substeps in the exact order shown; failure to do so might cause an outage or even data loss.

The following procedure sets node3 to boot from the root aggregate of node1:

a. Check the RAID, plex, and checksum information for the node1 aggregate:

```
aggr status -r
```

b. Check the status of the node1 aggregate:

```
aggr status
```

c. Bring the node1 aggregate online, if necessary:

```
aggr_online root_aggr_from_node1
```

d. Prevent the node3 from booting from its original root aggregate:


```
aggr offline root_aggr_on_node3
```

- e. Set the node1 root aggregate as the new root aggregate for node3:

```
aggr options aggr_from_node1 root
```

- f. Verify that the root aggregate of node3 is offline and the root aggregate for the disks brought over from node1 is online and set to root:

```
aggr status
```



Failing to perform the previous substep might cause node3 to boot from the internal root aggregate, or it might cause the system to assume a new cluster configuration exists or prompt you to identify one.

The following shows an example of the command output:

```
-----  
      Aggr State      Status      Options  
aggr0_nst_fas8080_15 online  raid_dp, aggr  root, nosnap=on  
                        fast zeroed  
                        64-bit  
  
      aggr0 offline      raid_dp, aggr  diskroot  
                        fast zeroed  
                        64-bit  
-----
```

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

```
ha-config show
```

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 programmable ROM (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
```

```
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
```

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

27. 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?
```

28. Enter `y` at the prompt to confirm that you want to destroy the local mailboxes.

29. Exit maintenance mode:

```
halt
```

The system stops at the boot environment prompt.

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

```
date
```

31. On node3, check the date at the boot environment prompt:

```
show date
```

32. If necessary, set the date on node3:

```
set date mm/dd/yyyy
```

33. On node3, check the time at the boot environment prompt:

```
show time
```

34. If necessary, set the time on node3:

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

35. Verify the partner system ID is set correctly as noted in [Step 21](#) under `-p` switch:

```
printenv partner-sysid
```

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

```
setenv partner-sysid node2_sysid
```

Save the settings:

```
saveenv
```

37. Access the boot menu at the boot environment prompt:

```
boot_ontap menu
```

38. 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?:
```

39. Enter `y` at the prompt.

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



The system might reboot twice before displaying the mismatch warning.

40. Confirm the mismatch as shown in the following 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.

41. Log in to node3.

Set the FC or UTA/UTA2 configuration on node3

If node3 has onboard FC ports, onboard unified target adapter (UTA/UTA2) ports, or a UTA/UTA2 card, you must configure the settings before completing the rest of the procedure.

About this task

You might need to complete [Configure FC ports on node3](#) or [Check and configure UTA/UTA2 ports on node3](#), or both sections.



NetApp marketing materials might use the term "UTA2" to refer to CNA adapters and ports. However, the CLI uses the term "CNA".

If node3 does not have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card, and you are upgrading a system with storage disks, you can skip to [Map ports from node1 to node3](#).

Configure FC ports on node3

If node3 has FC ports, either onboard or on an add-on FC adapter, you must set port configurations on the node before you bring it into service because the ports are not preconfigured when the systems are shipped. If you don't configure the ports, 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 [Prepare the nodes for upgrade](#).

About this task

You can skip this section if your system does not have FC configurations. If your system has onboard UTA/UTA2 ports or a UTA/UTA2 card, you configure them in [Check and configure UTA/UTA2 ports on node3](#).



Enter the commands in this section at the cluster prompt.

Steps

- 1. Display information about all FC and converged network adapters on the system.

```
system node hardware unified-connect show
```

- 2. Compare the FC settings of node3 with the settings that you captured earlier from node1.
- 3. 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 9 .
Different from the ones that you captured on node1	Go to Step 4 .

- 4. Modify the FC ports on node3 as needed by entering one of the following commands:

- To program target ports:

```
system node hardware unified-connect modify -type \|-t target -adapter
port_name
```

- To program initiator ports:

```
system node hardware unified-connect modify -type \|-t initiator -adapter
port_name
```

-t is the FC4 type: target or initiator.

- 5. Verify the new settings by entering the following command and examining the output:

```
system node hardware unified-connect show
```

- 6. Exit Maintenance mode:

```
halt
```

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

8. Boot node3 at the boot environment prompt:

```
boot_ontap
```

9. Take one of the following actions:

- If node3 has a UTA/UTA2 card or UTA/UTA2 onboard ports, go to [Check and configure UTA/UTA2 ports on node3](#).
- If node3 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports, skip [Check and configure UTA/UTA2 ports on node3](#) and go to [Map ports from node1 to node3](#).

Check and configure UTA/UTA2 ports on node3

If node3 has onboard UTA/UTA2 ports or a UTA/UTA2 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 UTA/UTA2 ports.

About this task

If you want to use a Unified Target Adapter (UTA/UTA2) port for FC, you must first verify how the port is configured.



NetApp marketing materials might use the term UTA2 to refer to CNA adapters and ports. However, the CLI uses the term CNA.

You can use the `ucadmin show` command to verify the current port configuration:

```
*> ucadmin show
```

Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
0e	fc	target	-	initiator	offline
0f	fc	target	-	initiator	offline
0g	fc	target	-	initiator	offline
0h	fc	target	-	initiator	offline
1a	fc	target	-	-	online
1b	fc	target	-	-	online

6 entries were displayed.

UTA/UTA2 ports can be configured into native FC mode or UTA/UTA2 mode. FC mode supports FC initiator and FC target; UTA/UTA2 mode enables concurrent NIC and FCoE traffic sharing the same 10GbE SFP+ interface and supports FC targets.

UTA/UTA2 ports might be found on an adapter or on the controller, and have the following configurations, but you should check the configuration of the UTA/UTA2 ports on the node3 and change it, if necessary:

- UTA/UTA2 cards ordered when the controller is ordered are configured before shipment to have the personality you request.

- UTA/UTA2 cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard UTA/UTA2 ports on new controllers are configured before shipment to have the personality you request.



Enter the commands in this section at the cluster prompt unless directed to enter Maintenance mode.

Steps

1. Check the current port configuration by entering the following command on node3:

```
system node hardware unified-connect show
```

The system displays output similar to the following example:

```
cluster1::> system node hardware unified-connect show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
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.

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 to determine whether the UTA/UTA2 ports have the personality you want.
4. Take one of the following actions:

If the UTA/UTA2 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 13 and go to Step 14 .

5. If the system has storage disks and is running clustered Data ONTAP 8.3, boot node3 and enter maintenance mode:

```
boot_ontap maint
```

6. Verify the settings:

```
ucadmin show
```

7. Take one of the following actions:

If you are configuring...	Then...
Ports on a UTA/UTA2 card	Go to Step 8 .
Onboard UTA/UTA2 ports	Skip Step 8 and go to Step 9 .

8. If the adapter is in initiator mode, and if the UTA/UTA2 port is online, take the UTA/UTA2 port offline:

```
storage disable adapter adapter_name
```

Adapters in target mode are automatically offline in maintenance mode.

9. 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` is the FC4 type, `target` or `initiator`.



You must use the FC initiator for tape drives and MetroCluster configurations. You must use the FC target for SAN clients.

10. Stop the system:

```
halt
```

The system stops at the boot environment prompt.

11. Enter the following command:

```
boot_ontap
```

12. Verify the settings:

```
system node hardware unified-connect show
```

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

```
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.

13. Place any target ports online by entering the following command, once for each port:

```
network fcp adapter modify -node node_name -adapter adapter_name -state up
```

14. Cable the port.

Map ports from node1 to node3

You must 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*. (Go to [References](#) to link to the *Hardware Universe*). You use the information later in this section and in [Map ports from node2 to node4](#).

The software configuration of node3 must match the physical connectivity of node3, and network 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 might 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 domain	Node3 ports	Node3 ports	Node3 broadcast domains
Cluster 1						
Cluster 2						
Cluster 3						
Cluster 4						
Cluster 5						
Cluster 6						
Node management						
Cluster management						
Data 1						
Data 2						
Data 3						
Data 4						
SAN						
Intercluster port						

Refer to [Record node1 information](#) for the steps to obtains 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 [Record node1 information](#).
3. Follow these steps to verify if the setup is a two-node switchless cluster:
 - a. Set the privilege level to advanced:

```
cluster::> set -privilege advanced
```

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

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

```
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 [Install and boot node3](#) 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-name -port port-name -fields broadcast-domain
```

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 node-name -port port-name -ip-space Cluster -mtu
9000
```

This example adds Cluster port "e1b" on node3:

```
network port modify -node node3 -port e1b -ip-space Cluster -mtu 9000
```



For a 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 sync-destination SVM and see errors similar to, but not restricted to the following message`:

```
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:

```
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
```

The following example 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
```

5. Adjust the broadcast domains hosting your cluster LIFs and node-management and/or cluster-management LIFs. Confirm 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 might need to migrate and modify the LIFs as follows:

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 Vserver-name -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 5](#).

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 node1 that no longer exist on node3, follow these steps to delete them:

- a. Access the advanced privilege level on either node:

```
set -privilege advanced
```

- b. Delete the ports:

```
network port delete -node node-name -port port-name
```

- c. 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
```

The following example 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-domainwide -failover-group fg1
```

Go to [References](#) to link to *Network Management* or the *ONTAP 9 Commands: Manual Page Reference* for more information.

10. Verify the changes on node3:

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

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

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

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

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

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

then up:

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

Repeat Step 11 to verify that the cluster LIF is now listening on port 7700.

Verify the node3 installation

After you install and boot node3, you must verify that it is installed correctly, that it is part of the cluster, and that it can communicate with node2.

Steps

1. At the system prompt, log in to node3. Then, verify that node3 is both part of the same cluster as node2 and healthy:

```
cluster show
```

2. Verify that node3 can communicate with node2 and that all LIFs are up:

```
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 Moving NAS data LIFs owned by node1 from node2 to node3 and verifying SAN LIFs on node3 .
Not in a SAN environment	Skip Step 4 and go to Moving NAS data LIFs owned by node1 from node2 to node3 and verifying SAN LIFs on node3 .

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.*
```

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 ...
```

Move NAS data LIFs owned by node1 from node2 to node3 and verify SAN LIFs on node3

After you verify the node3 installation and before you relocate aggregates from node2 to node3, you must move the NAS 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 NAS 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. If the cluster is configured for SAN LIFs, record the SAN LIFs adapter and switch-port configuration information in this [worksheet](#) for use later in the procedure.
 - a. List the SAN LIFs on node2 and examine the output:

```
network interface show -data-protocol fc*
```

The system returns output similar to the following example:

```

cluster1::> net int show -data-protocol fc*
(network interface show)

```

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			
-----	-----	-----	-----	-----
-----	-----			
svm2_cluster1	lif_svm2_cluster1_340	up/up	20:02:00:50:56:b0:39:99	cluster1-01
1b	true			
	lif_svm2_cluster1_398	up/up	20:03:00:50:56:b0:39:99	cluster1-02
1a	true			
	lif_svm2_cluster1_691	up/up	20:01:00:50:56:b0:39:99	cluster1-01
1a	true			
	lif_svm2_cluster1_925	up/up	20:04:00:50:56:b0:39:99	cluster1-02
1b	true			

4 entries were displayed.

b. List the existing configurations and examine the output:

```
fcip adapter show -fields switch-port,fc-wwpn
```

The system returns output similar to the following example:

```
cluster1::> fcp adapter show -fields switch-port,fc-wwpn
(network fcp adapter show)
node          adapter  fc-wwpn                      switch-port
-----
cluster1-01   0a          50:0a:09:82:9c:13:38:00      ACME Switch:0
cluster1-01   0b          50:0a:09:82:9c:13:38:01      ACME Switch:1
cluster1-01   0c          50:0a:09:82:9c:13:38:02      ACME Switch:2
cluster1-01   0d          50:0a:09:82:9c:13:38:03      ACME Switch:3
cluster1-01   0e          50:0a:09:82:9c:13:38:04      ACME Switch:4
cluster1-01   0f          50:0a:09:82:9c:13:38:05      ACME Switch:5
cluster1-01   1a          50:0a:09:82:9c:13:38:06      ACME Switch:6
cluster1-01   1b          50:0a:09:82:9c:13:38:07      ACME Switch:7
cluster1-02   0a          50:0a:09:82:9c:6c:36:00      ACME Switch:0
cluster1-02   0b          50:0a:09:82:9c:6c:36:01      ACME Switch:1
cluster1-02   0c          50:0a:09:82:9c:6c:36:02      ACME Switch:2
cluster1-02   0d          50:0a:09:82:9c:6c:36:03      ACME Switch:3
cluster1-02   0e          50:0a:09:82:9c:6c:36:04      ACME Switch:4
cluster1-02   0f          50:0a:09:82:9c:6c:36:05      ACME Switch:5
cluster1-02   1a          50:0a:09:82:9c:6c:36:06      ACME Switch:6
cluster1-02   1b          50:0a:09:82:9c:6c:36:07      ACME Switch:7
16 entries were displayed
```

3. Take one of the following actions:

If node1...	Then...
Had interface groups or VLANs configured	Go to Step 4 .
Did not have interface groups or VLANs configured	Skip Step 4 and go to Step 5 .

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

- 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
```

- 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
```

- 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
```

5. Take one of the following actions:

If the cluster is configured for...	Then...
NAS	Complete Step 6 and Step 7 , skip Step 8, and complete Step 9 through Step 12 .
SAN	Disable all the SAN LIFs on the node to take them down for the upgrade: <pre>network interface modify -vserver vservice_name -lif LIF_name -home-node node_to_upgrade -home-port netport ifgrp -status-admin down</pre>

6. If you have data ports that are not the same on your platforms, then add the ports to the broadcast domain:

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

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
```

7. Migrate each NAS 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
```

8. Make sure that the data migration is persistent:

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

9. Confirm that the SAN LIFs are on the correct ports on node3:

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

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

The system returns output similar to the following example:

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

	Logical	Status	Network	Current
Current Is				
Vserver	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. Verify that the new and adapter and switch-port configurations are correct by comparing the output from the `fc adapter show` command with the configuration information that you recorded in the worksheet in [Step 2](#).

List the new SAN LIF configurations on node3:

```
fc adapter show -fields switch-port,fc-wwpn
```

The system returns output similar to the following example:

```
cluster1::> fcp adapter show -fields switch-port,fc-wwpn
(network fcp adapter show)
node          adapter fc-wwpn          switch-port
-----
cluster1-01 0a      50:0a:09:82:9c:13:38:00 ACME Switch:0
cluster1-01 0b      50:0a:09:82:9c:13:38:01 ACME Switch:1
cluster1-01 0c      50:0a:09:82:9c:13:38:02 ACME Switch:2
cluster1-01 0d      50:0a:09:82:9c:13:38:03 ACME Switch:3
cluster1-01 0e      50:0a:09:82:9c:13:38:04 ACME Switch:4
cluster1-01 0f      50:0a:09:82:9c:13:38:05 ACME Switch:5
cluster1-01 1a      50:0a:09:82:9c:13:38:06 ACME Switch:6
cluster1-01 1b      50:0a:09:82:9c:13:38:07 ACME Switch:7
cluster1-02 0a      50:0a:09:82:9c:6c:36:00 ACME Switch:0
cluster1-02 0b      50:0a:09:82:9c:6c:36:01 ACME Switch:1
cluster1-02 0c      50:0a:09:82:9c:6c:36:02 ACME Switch:2
cluster1-02 0d      50:0a:09:82:9c:6c:36:03 ACME Switch:3
cluster1-02 0e      50:0a:09:82:9c:6c:36:04 ACME Switch:4
cluster1-02 0f      50:0a:09:82:9c:6c:36:05 ACME Switch:5
cluster1-02 1a      50:0a:09:82:9c:6c:36:06 ACME Switch:6
cluster1-02 1b      50:0a:09:82:9c:6c:36:07 ACME Switch:7
16 entries were displayed
```



If a SAN LIF in the new configuration is not on an adapter that is still attached to the same switch-port, it might cause a system outage when you reboot the node.

- c. 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":

```
network interface modify -vserver vservice_name -lif LIF_name -status
-admin down
```

- ii. Remove the LIF from the port set:

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

- iii. Enter one of the following commands:

- Move a single LIF:

```
network interface modify -vserver vservice_name -lif LIF_name -home
-port new_home_port
```

- Move all the LIFs on a single nonexistent or incorrect port to a new port:

```
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
```



You must move SAN LIFs to a port that has the same link speed as the original port.

10. Modify the status of all LIFs to "up" so the LIFs can accept and send traffic on the node:

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

11. Enter the following command on either node 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
```

12. 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
```

13. Send a post-upgrade AutoSupport message to NetApp for node1:

```
system node autosupport invoke -node node3 -type all -message "node1  
successfully upgraded from platform_old to platform_new"
```

Worksheet: Information to record before moving NAS data LIFs to node3

To help verify that you have the correct configuration after moving SAN LIFs from node2 to node3, you can use the following worksheet to record the adapter and switch-port information for each LIF.

Record the LIF adapter information from the `network interface show -data-protocol fc*` command output and the switch-port information from the `fcport adapter show -fields switch-port,fc-wwpn` command output for node2.

After you complete the migration to node3, record the LIF adapter and switch-port information for the LIFs on node3 and verify that each LIF is still connected to the same switch-port.

Node2			Node3		
LIF	adapter	switch-port	LIF	adapter	switch-port

Node2			Node3		

Relocate non-root aggregates from node2 to node3

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



During this procedure, don't relocate aggregates from node3 to node2. Doing so results in aggregates being taken offline and a data outage for the aggregates that are relocated.

Steps

1. Verify that the partner system ID is set correctly on node3:

- a. Enter the advanced privilege level:

```
set -privilege advanced
```

- b. Show the partner system ID on node3:

```
ha interconnect config show -node <node3-node1>
```

The system displays output similar to the following example:

Show example

```
cluster::*> ha interconnect config show -node <node>
(system ha interconnect config show)

                        Node: node3-node1
      Interconnect Type: RoCE
        Local System ID: <node3-system-id>
        Partner System ID: <node2-system-id>
    Connection Initiator: local
                Interface: external

Port   IP Address
----   -
e4a-17 0.0.0.0
e4b-18 0.0.0.0
```

2. If "Partner System ID" is incorrect for node3:

- a. Halt node3:

```
halt
```

- b. At the LOADER prompt, set the correct "partner-sysid" value.

The node3 "partner-sysid" is the system ID of node2, which you can find in the `ha interconnect config show` output in [Step 1](#).

- c. Save the settings:

```
saveenv
```

- d. At the LOADER prompt, boot node3 into the boot menu:

```
boot_ontap menu
```

- e. Log in to node3.

3. 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>"
```

4. 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.

5. Relocate the non-root aggregates:

- a. Set the privilege level to advanced:

```
set -privilege advanced
```

- b. List the aggregates that are owned by node2:

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

- c. Start aggregate relocation:

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



The command locates only non-root aggregates.

- d. 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 doesn't relocate any offline or restricted aggregates.

- e. Return to the admin privilege level:

```
set -privilege admin
```

6. Verify the relocation status of node2:

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


The output displays "Done" for an aggregate after it has been relocated.



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

7. Take one of the following actions:

If relocation of...	Then...
All aggregates was successful	Go to Step 8 .

If relocation of...	Then...
Any aggregates failed, or was vetoed	<p>a. Display a detailed status message:</p> <pre>storage aggregate show -instance</pre> <p>You can also check the EMS logs to see the corrective action that is needed.</p> <div style="display: flex; align-items: center; margin-top: 10px;">  <div> <p>The event log show command lists any errors that have occurred.</p> </div> </div> <p>b. Perform the corrective action.</p> <p>c. Set the privilege level to advanced:</p> <pre>set -privilege advanced</pre> <p>d. Relocate any failed or vetoed aggregates:</p> <pre>storage aggregate relocation start -node <node2> -destination <node3> -aggregate-list * -ndo-controllerupgrade true</pre> <p>e. When prompted, enter y.</p> <p>f. Return to the admin privilege level:</p> <pre>set -privilege admin</pre> <p>If necessary, you can force the relocation by using one of the following methods:</p> <ul style="list-style-type: none"> • By overriding veto checks: <pre>storage aggregate relocation start -override -vetoes true -ndo-controller-upgrade</pre> • By overriding destination checks: <pre>storage aggregate relocation start -override -destination-checks true -ndocontroller-upgrade</pre> <p>For more information about the storage aggregate relocation commands, go to References to link to <i>Disk and aggregate management with the CLI</i> and the <i>ONTAP 9 Commands: Manual Page Reference</i>.</p>

8. 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 <aggregate_name>
```

9. 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 <Vserver-name> -volume <volume-name>
```

10. Verify that node2 doesn't 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.

Move NAS data LIFs owned by node2 to node3

After you relocate the aggregates from node2 to node3, you need to move the NAS data LIFs owned by node2 to 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 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

1. List all the NAS 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
```

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...	Then...
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 .

- Take the following steps to migrate NAS data LIFs hosted on interface groups and VLANs on node2:
 - 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 Vserver_name -lif LIF_name -destination
-node node3 -destination-port netport|ifgrp
```

- 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 Vserver_name -lif LIF_name -home-node
node3 -homeport netport|ifgrp
```

- 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 Vserver_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 Vserver_name -lif LIF_name -home-node
node3 -homeport 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, add the ports to the broadcast domain:

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

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
```

6. Migrate each NAS 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. Verify that NAS 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 Vserver_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:

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

- b. Enter the following command and examine its output to determine 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: MAC_address
      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
```

Stage 4. Record information and retire node2

Record node2 information

Before you can shut down and retire node2, you must 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,nodemgmt,cluster-mgmt
```

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

```

Is	Logical	Status	Network	Current	Current
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

4 entries were displayed.



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 is displayed in the example output of [Step 1](#) in *Record node1 port information*.

2. Capture the information in the output to use in the section [Map 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.

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

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

(Mbps)					Speed	
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper
-----	-----	-----	-----	-----	-----	-----
node2						
	e0M	Default	IP_address	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
```

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 [Map ports from node2 to node4](#).

- Take one of the following actions:

If you...	Then...
Recorded NVRAM System ID number in Prepare the nodes for upgrade	Go to Retire node2 .
Did not record the NVRAM System ID number in Prepare the nodes for upgrade	Complete Step 9 and Step 10 and then go to the next section, Retire node2 .

- Display the attributes of node2:

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

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

- Record the NVRAM System ID to use in the section [Install and boot node4](#).

Retire node2

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

Steps

- Take one of the following actions:

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

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

```
set -privilege advanced
```

- 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
```

4. Check if node2 currently holds epsilon by entering the following command and examining its output:

```
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.



If you are upgrading an HA pair in a cluster with multiple HA pairs, you must move epsilon to the node of an HA pair that isn't undergoing a controller upgrade. For example, if you are upgrading nodeA/nodeB in a cluster with the HA pair configuration nodeA/nodeB and nodeC/nodeD, you must move epsilon to nodeC or nodeD.

5. If node2 holds epsilon, mark epsilon as false on the node so that it can be transferred to node3:

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

6. Transfer epsilon to node3 by marking epsilon true on node3:

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

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

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

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

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

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

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



```
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. Return to the admin level:

```
set -privilege admin
```

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

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

11. After node2 shuts down completely, remove it from the chassis or the rack. You can decommission node2 after the upgrade is completed. See [Decommission the old system](#).

Stage 5. Install and boot node4

Install and boot node4

You must 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 must netboot node4 if the ONTAP version on node4 is different to the ONTAP version on node2. After you install node4, boot it 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 by following the instructions in [Prepare for netboot](#)

However, you don't need to netboot node4 if the ONTAP version on node4 is the same or later than the ONTAP version on node2.



- For an AFF A800 or AFF C800 controller upgrade, you must ensure that all drives in the chassis are firmly seated against the midplane before removing node2. For more information, see [Replace the AFF A800 or AFF C800 controller modules](#).
- If you're upgrading a system with storage disks, you must complete this entire section and then proceed to the section [Set the FC or UTA/UTA2 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 Steps 2 and 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.

The following references help you make proper cable connections. Go to [References](#) to link to them.

- *Installation and Setup Instructions* for the node4 platform
- The appropriate disk shelf procedure
- The *HA pair management* documentation

Cable the following connections:

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



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.



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:
 - 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 finish.



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

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

```
halt
```

The system stops at the boot environment prompt.

9. Configure node4 for ONTAP:

```
set-defaults
```

10. If you have NetApp Storage Encryption (NSE) drives installed, perform the following steps:



If you have not already done so earlier in the procedure, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use.

- a. Set `bootarg.storageencryption.support` to `true` or `false`:

If the following drives are in use...	Then...
NSE drives that conform to FIPS 140-2 Level 2 self-encryption requirements	<code>setenv bootarg.storageencryption.support true</code>
NetApp non-FIPS SEDs	<code>setenv bootarg.storageencryption.support false</code>



You cannot mix FIPS drives with other types of drives on the same node or HA pair.

You can mix SEDs with non-encrypting drives on the same node or HA pair.

- b. Contact NetApp Support for assistance with restoring the onboard key management information.

11. If the version of ONTAP installed on node4 is the same or later than the version of ONTAP 9 installed on node2, enter the following command:

```
boot_ontap menu
```


12. Take one of the following actions:

If the system you are upgrading...	Then...
Does not have the correct or current ONTAP version on node4	Go to Step 13 .
Has the correct or current version of ONTAP on node4	Go to Step 18 .

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



You must use the management port and IP address as the netboot connection. Do not use a data LIF IP address 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: <code>ifconfig e0M -auto</code>
Not running	<p>Manually configure the connection by entering the following command at the boot environment prompt: <code>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></code></p> <p><i>filer_addr</i> is the IP address of the storage system (mandatory). <i>netmask</i> is the network mask of the storage system (mandatory). <i>gateway</i> is the gateway for the storage system (mandatory). <i>dns_addr</i> is the IP address of a name server on your network (optional). <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> <div>  <p>Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p> </div>

14. Perform netboot on node4:

For...	Then...
FAS/AFF8000 series systems	<code>netboot http://<web_server_ip/path_to_webaccessible_directory> /netboot/kernel</code>
All other systems	<code>netboot http://<web_server_ip/path_to_webaccessible_directory/ ontap_version>_image.tgz</code>

The `<path_to_the_web-accessible_directory>` should lead to where you downloaded the `<ontap_version>_image.tgz` in [Step 1](#) in the section *Prepare for netboot*.



Do not interrupt the boot.

15. 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.



Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the incorrect image might install. This issue applies to all releases of ONTAP. The netboot procedure combined with option (7) Install new software wipes the boot media and places the same ONTAP version on both image partitions.

16. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL:

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

17. 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}
```

- 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.

18. Select maintenance mode 5 from the boot menu and enter `y` when you are prompted to continue with the boot.
19. Before continuing, go to [Set the FC or UTA/UTA2 configuration on node4](#) to make any necessary changes to the FC or UTA/UTA2 ports on the node. Make the changes recommended in those sections, reboot the node, and go into Maintenance mode.
20. Enter the following command and examine the output to find the system ID of node4:

```
disk show -a
```

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
-----
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)
.....
```

21. Reassign node2's spares, disks belonging to the root, and any non-root aggregates that were not relocated to node3 earlier in section [Relocate non-root aggregates from node2 to node3](#):



If you have shared disks, hybrid aggregates, or both on your system, you must use the correct `disk reassign` command from the following table.

Disk type...	Run the command...
With shared disks	<pre>disk reassign -s node2_sysid -d node4_sysid -p node3_sysid</pre>
Without shared	<pre>disks disk reassign -s node2_sysid -d node4_sysid</pre>

For the `<node2_sysid>` value, use the information captured in [Step 10](#) of the *Record node2 information* section. For `node4_sysid`, use the information captured in [Step 23](#).



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
```

Enter `n` when asked to abort disk reassignment.

When you are asked to abort disk reassignment, you must answer a series of prompts as shown in the following steps:

- a. 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
```

- b. Enter `y` to continue.

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
```

- c. Enter `y` to allow disk ownership to be updated.

22. If you are upgrading from a system with external disks to a system that supports internal and external disks (A800 systems, for example), set `node4` as root to confirm that it boots from the root aggregate of `node2`.



Warning: You must perform the following substeps in the exact order shown; failure to do so might cause an outage or even data loss.

The following procedure sets `node4` to boot from the root aggregate of `node2`:

- a. Check the RAID, plex, and checksum information for the `node2` aggregate:

```
aggr status -r
```

- b. Check the overall status of the `node2` aggregate:

```
aggr status
```

- c. If necessary, bring the `node2` aggregate online:

```
aggr_online root_aggr_from_node2
```

d. Prevent the node4 from booting from its original root aggregate:

```
aggr offline root_aggr_on_node4
```

e. Set the node2 root aggregate as the new root aggregate for node4:

```
aggr options aggr_from_node2 root
```

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

```
ha-config show
```

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
```

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

If you have a MetroCluster configuration, use the following commands to correct the configuration:

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

```
ha-config modify chassis mcc.
```

24. Destroy the mailboxes on node4:

```
mailbox destroy local
```

25. Exit Maintenance mode:

```
halt
```

The system stops at the boot environment prompt.

26. On node3, check the system date, time, and time zone:

```
date
```

27. On node4, check the date at the boot environment prompt:

```
show date
```


28. If necessary, set the date on node4:

```
set date mm/dd/yyyy
```

29. On node4, check the time at the boot environment prompt:

```
show time
```

30. If necessary, set the time on node4:

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

31. Verify the partner system ID is set correctly as noted in [Step 19](#) under option.

```
printenv partner-sysid
```

32. If necessary, set the partner system ID on node4:

```
setenv partner-sysid node3_sysid
```

a. Save the settings:

```
saveenv
```

33. Enter the boot menu at the boot environment prompt:

```
boot_ontap menu
```

34. 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?:
```

35. Enter **y** at the prompt.

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



The system might reboot twice before displaying the mismatch warning.

36. Confirm the mismatch.

The node might complete one round of rebooting before booting normally.

37. Log in to node4.

Set the FC or UTA/UTA2 configuration on node4

If node4 has onboard FC ports, onboard unified target adapter (UTA/UTA2) ports, or a UTA/UTA2 card, you must configure the settings before completing the rest of the procedure.

About this task

You might need to complete [Configure FC ports on node4](#) or [Check and configure UTA/UTA2 ports on node4](#), or both sections.

If node4 does not have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card, and you are upgrading a system with storage disks, you can skip to [Map ports from node2 to node4](#).

Configure FC ports on node4

If node4 has FC ports, either 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 [Prepare the nodes for upgrade](#).

About this task

You can skip this section if your system does not have FC configurations. If your system has onboard UTA/UTA2 ports or a UTA/UTA2 adapter, you configure them in [Check and configure UTA/UTA2 ports on node4](#).



Enter the commands in this section at the cluster prompt.

Steps

1. Display information about all FC and converged network adapters on the system:

```
system node hardware unified-connect show
```

2. Compare the FC settings on the new nodes with the settings that you captured earlier from the original node.

3. Modify the FC ports on node4 as needed:

- To program target ports:

```
system node hardware unified-connect modify -type \|-t target -adapter  
port_name
```

- To program initiator ports:

```
system node unified-connect modify type \|-t initiator -adapter port_name
```

-type is the FC4 type, target or initiator.

4. Verify the new settings by entering the following command and examining the output:

```
system node unified-connect show
```

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 the original nodes	Go to Step 9 .
Different from the ones that you captured on the original nodes	Go to Step 6 .

6. Exit Maintenance mode:

```
halt
```

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

8. Boot node4 by entering the following command at the boot environment prompt:

```
boot_ontap
```

9. Take one of the following actions:

- Go to [Check and configure UTA/UTA2 ports on node4](#) if node4 has a UTA/UTA2A card or UTA/UTA2 onboard ports.
- Skip the section and go to [Map ports from node2 to node4](#) if node4 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports.

Check and configure UTA/UTA2 ports on node4

If node4 has onboard UTA/UTA2 ports or a UTA/UTA2A 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 UTA/UTA2 ports.

About this task

UTA/UTA2 ports can be configured into native FC mode or UTA/UTA2A mode. FC mode supports FC initiator and FC target; UTA/UTA2 mode enables concurrent NIC and FCoE traffic to share the same 10GbE SFP+ interface and supports FC target.



NetApp marketing materials might use the term UTA2 to refer to CNA adapters and ports. However, the CLI uses the term CNA.

UTA/UTA2 ports might be on an adapter or on the controller with the following configurations:

- UTA/UTA2 cards ordered at the same time as the controller are configured before shipment to have the personality you requested.
- UTA/UTA2 cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard UTA/UTA2 ports on new controllers are configured (before shipment) to have the personality you requested.

However, you can check the configuration of the UTA/UTA2 ports on node4 and change it, if necessary.



Enter the commands in this section at the cluster prompt unless directed to enter Maintenance mode. If you have a MetroCluster FC system, you must be in Maintenance mode to configure UTA/UTA2 ports.

Steps

1. Check how the ports are currently configured on node4:

```
system node hardware unified-connect show
```

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 UTA/UTA2 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 to Step 13 .

5. If the system has storage disks and is running Data ONTAP 8.3, boot node4 and enter maintenance mode:

```
boot_ontap maint
```

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

```
ucadmin show
```

7. Take one of the following actions:

If you are configuring...	Then...
Ports on a UTA/UTA2A card	Go to Step 8 .
Onboard UTA/UTA2 ports	Skip Step 8 and go to Step 9 .

8. If the adapter is in initiator mode, and if the UTA/UTA2 port is online, take the UTA/UTA2 port offline:

```
storage disable adapter adapter_name
```

Adapters in target mode are automatically offline in Maintenance mode.

9. 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 UTA.

- `-t` is the FC4 type: target or initiator.



You must use FC initiator for tape drives and the FC target for SAN clients.

- If the system has storage disks, enter the following command:

```
halt
```

The system stops at the boot environment prompt.

- Enter the following command:

```
boot_ontap
```

- If the system has storage disks, enter the following command:

```
system node hardware unified-connect show
```

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.

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

```
network fcp adapter modify -node node_name -adapter adapter_name -state up
```

- Cable the port.

Map ports from node2 to node4

You must 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

You must already have information about the ports on the new nodes, to access this information refer to [References](#) to link to the *Hardware Universe*. You use the information later in this section.

The software configuration of node4 must match the physical connectivity of node4, 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, 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

- 1. 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						
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 [Record 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:

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

```
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 [Install and boot 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`
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
```

This example adds Cluster port "e1b" on node4:

```
network port modify -node node4 -port e1b -ip-space Cluster -mtu 9000
```



For a 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 sync-destination SVM and see errors similar to, but not restricted to, the following:

```
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
```

This command removes port "e0d" on node4:

```
network port broadcast-domain remove-ports -ipSPACE 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. Confirm 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 vserver_name -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 5](#).
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
```

- b. To delete the ports:

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

- c. 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
```

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
```

For more information, refer to [References](#) to link to *Network Management* or the *ONTAP 9 Commands: Manual Page Reference*, and go to *Configuring failover settings on a LIF*.

10. Verify the changes on node4:

```
network port show -node node4
```

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

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

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

```

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

```

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

```

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

```

Repeat Step 11 to verify that the cluster LIF is now listening on port 7700.

Verify the node4 installation

After you install and boot node4, you must verify that it is installed correctly, that it is part of the cluster, and that it can communicate with node3.

Steps

- At the system prompt, log in to node4.
- Verify that node4 is both part of the same cluster as node3 and healthy:

```
cluster show
```

- Verify that node4 can communicate with node3 and that all LIFs are up:

```
network interface show -curr-node node4
```

- Take one of the following actions:

If node4 is...	Then...
In a chassis separate from node3	<p>Connect the HA interconnect between the nodes by completing the following steps:</p> <ol style="list-style-type: none"> Connect the top interconnect port of node3 to the top interconnect port of node4. Connect the bottom interconnect port of node3 to the bottom interconnect port of node4. Go to Step 5.

If node4 is...	Then...
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 Move NAS data LIFs owned by node2 from node3 to node4 and verify SAN LIFs on node4 .
Not in a SAN environment	Skip Step 6 go to the section Move NAS data LIFs owned by node2 from node3 to node4 and verify SAN LIFs on node4 .

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.*
```

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 ...
```

Move NAS data LIFs owned by node2 from node3 to node4 and verify SAN LIFs on node4

After you verify the node4 installation and before you relocate node2 aggregates from node3 to node4, you must move the NAS 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 NAS 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. If the cluster is configured for SAN LIFs, record the SAN LIFs and existing configuration information in this [worksheet](#) for use later in the procedure.

- a. List the SAN LIFs on node3 and examine the output:

```
network interface show -data-protocol fc*
```

The system returns output similar to the following example:

```
cluster1::> net int show -data-protocol fc*
(network interface show)
      Logical      Status      Network      Current
Current Is
Vserver      Interface      Admin/Oper      Address/Mask      Node
Port      Home
-----
-----
svm2_cluster1
      lif_svm2_cluster1_340
                        up/up      20:02:00:50:56:b0:39:99
                        cluster1-01
1b      true
      lif_svm2_cluster1_398
                        up/up      20:03:00:50:56:b0:39:99
                        cluster1-02
1a      true
      lif_svm2_cluster1_691
                        up/up      20:01:00:50:56:b0:39:99
                        cluster1-01
1a      true
      lif_svm2_cluster1_925
                        up/up      20:04:00:50:56:b0:39:99
                        cluster1-02
1b      true
4 entries were displayed.
```

- b. List the existing configurations and examine the output:

```
fcv adapter show -fields switch-port,fc-wwpn
```

The system returns output similar to the following example:

```
cluster1::> fcp adapter show -fields switch-port,fc-wwpn
(network fcp adapter show)
node          adapter  fc-wwpn                      switch-port
-----
cluster1-01   0a          50:0a:09:82:9c:13:38:00     ACME Switch:0
cluster1-01   0b          50:0a:09:82:9c:13:38:01     ACME Switch:1
cluster1-01   0c          50:0a:09:82:9c:13:38:02     ACME Switch:2
cluster1-01   0d          50:0a:09:82:9c:13:38:03     ACME Switch:3
cluster1-01   0e          50:0a:09:82:9c:13:38:04     ACME Switch:4
cluster1-01   0f          50:0a:09:82:9c:13:38:05     ACME Switch:5
cluster1-01   1a          50:0a:09:82:9c:13:38:06     ACME Switch:6
cluster1-01   1b          50:0a:09:82:9c:13:38:07     ACME Switch:7
cluster1-02   0a          50:0a:09:82:9c:6c:36:00     ACME Switch:0
cluster1-02   0b          50:0a:09:82:9c:6c:36:01     ACME Switch:1
cluster1-02   0c          50:0a:09:82:9c:6c:36:02     ACME Switch:2
cluster1-02   0d          50:0a:09:82:9c:6c:36:03     ACME Switch:3
cluster1-02   0e          50:0a:09:82:9c:6c:36:04     ACME Switch:4
cluster1-02   0f          50:0a:09:82:9c:6c:36:05     ACME Switch:5
cluster1-02   1a          50:0a:09:82:9c:6c:36:06     ACME Switch:6
cluster1-02   1b          50:0a:09:82:9c:6c:36:07     ACME Switch:7
16 entries were displayed
```

3. Take one of the following actions:

If node2...	Description
Had interface groups or VLANs configured	Go to Step 4 .
Did not have interface groups or VLANs configured	Skip Step 4 and go to Step 5 .

4. Take the following steps to migrate any NAS data LIFs hosted on interface groups and VLANs that originally were on node2 from node3 to node4.

- 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 vservice_name -lif lif_name -destination
-node node4 -destination-port netport|ifgrp
```

- 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 vservice_name -lif datalif_name -home-node
node4 home-port netport|ifgrp
```

- 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 vservice_name -lif datalif_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 vservice_name -lif datalif_name -home-node
node4 home-port netport|ifgrp
```

5. Take one of the following actions:

If the cluster is configured for...	Then...
NAS	Complete Step 6 through Step 9 , skip Step 10 , and complete Step 11 through Step 14 .
SAN	Skip Step 6 through Step 9 , and complete Step 10 through Step 14 .
Both NAS and SAN	Complete Step 6 through Step 14 .

6. 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
```

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
```

7. Migrate each NAS 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
```

8. Make sure that the data migration is persistent:

```
network interface modify -vserver vservice_name -lif datalif_name -home-port
netport|ifgrp
```

9. 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
```

The following example shows the output of the `network port show` command with some LIFs up and others down:

```
cluster::> network port show
```

(Mbps)					Speed	
Node	Port	IPspace	Broadcast Domain	Link	MTU	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.						

10. 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:

- a. Enter the advanced privilege level by entering the following command:

```
set -privilege advanced
```

- b. Enter the following command, once for each old network port:

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

- c. Return to the admin level by entering the following command:

```
set -privilege admin
```

11. Confirm that the SAN LIFs are on the correct ports on node4 by completing the following substeps:

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

```
network interface show -data-protocol iscsi|fc -home-node node4
```

The system returns output similar to the following example:

```
cluster::> network interface show -data-protocol iscsi|fc -home-node
node4
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			
-----	-----	-----	-----	-----
-----	----			
vs0				
	a0a	up/down	10.63.0.53/24	node4
a0a	true			
	data1	up/up	10.63.0.50/18	node4
e0c	true			
	rads1	up/up	10.63.0.51/18	node4
e1a	true			
	rads2	up/down	10.63.0.52/24	node4
e1b	true			
vs1				
	lif1	up/up	172.17.176.120/24	node4
e0c	true			
	lif2	up/up	172.17.176.121/24	node4

- b. Verify that the new adapter and switch-port configurations are correct by comparing the output from the `fc -adapter show` command with the new configuration information that you recorded in the worksheet in [Step 2](#).

List the new SAN LIF configurations on node4:

```
fc -adapter show -fields switch-port,fc-wwpn
```

The system returns output similar to the following example:


```
cluster1::> fcp adapter show -fields switch-port,fc-wwpn
(network fcp adapter show)
node          adapter  fc-wwpn                switch-port
-----
cluster1-01   0a          50:0a:09:82:9c:13:38:00 ACME Switch:0
cluster1-01   0b          50:0a:09:82:9c:13:38:01 ACME Switch:1
cluster1-01   0c          50:0a:09:82:9c:13:38:02 ACME Switch:2
cluster1-01   0d          50:0a:09:82:9c:13:38:03 ACME Switch:3
cluster1-01   0e          50:0a:09:82:9c:13:38:04 ACME Switch:4
cluster1-01   0f          50:0a:09:82:9c:13:38:05 ACME Switch:5
cluster1-01   1a          50:0a:09:82:9c:13:38:06 ACME Switch:6
cluster1-01   1b          50:0a:09:82:9c:13:38:07 ACME Switch:7
cluster1-02   0a          50:0a:09:82:9c:6c:36:00 ACME Switch:0
cluster1-02   0b          50:0a:09:82:9c:6c:36:01 ACME Switch:1
cluster1-02   0c          50:0a:09:82:9c:6c:36:02 ACME Switch:2
cluster1-02   0d          50:0a:09:82:9c:6c:36:03 ACME Switch:3
cluster1-02   0e          50:0a:09:82:9c:6c:36:04 ACME Switch:4
cluster1-02   0f          50:0a:09:82:9c:6c:36:05 ACME Switch:5
cluster1-02   1a          50:0a:09:82:9c:6c:36:06 ACME Switch:6
cluster1-02   1b          50:0a:09:82:9c:6c:36:07 ACME Switch:7
16 entries were displayed
```



If a SAN LIF in the new configuration is not on an adapter that is still attached to the same switch-port, it might cause a system outage when you reboot the node.

- c. 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:

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

- ii. Remove the LIF from the port set:

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

- iii. Enter one of the following commands:

- Move a single LIF:

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

- Move all the LIFs on a single nonexistent or incorrect port to a new port:

```
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 vservice_name -portset portset_name -port-name
port_name
```



You must move SAN LIFs to a port that has the same link speed as the original port.

12. 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 vservers_name -home-port port_name -home-node
node4 lif lif_name -status-admin up
```

13. 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
```

14. 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 vsrvr name -lif lif name -status-admin up
```

Worksheet: Information to record before moving NAS data LIFs to node4

To help verify that you have the correct configuration after moving SAN LIFs from node3 to node4, you can use the following worksheet to record the `adapter` and `switch-port` information for each LIF.

Record the LIF adapter information from the network interface `show -data-protocol fc*` command output and the switch-port information from the fcp adapter `show -fields switch-port,fc-wwpn` command output for node3.

After you complete the migration to node4, record the LIF adapter and switch-port information for the LIFs on node4 and verify that each LIF is still connected to the same switch-port.

[illegible]

Relocate node2 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
```

- b. 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](#).

- c. 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.

- d. Return to the admin level:

```
set -privilege admin
```

- 3. Check the relocation status:

```
storage aggregate relocation show -node node3
```

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



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 .

If relocation of...	Then...
Any aggregates failed, or were 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:</p> <pre>set -privilege advanced</pre> <p>d. Relocate any failed or vetoed aggregates:</p> <pre>storage aggregate relocation start -node node3 destination node4 -aggregate-list aggr_name1, aggr_name2... ndo-controller-upgrade true</pre> <p>The aggregate list is the list of failed or vetoed aggregates.</p> <p>e. When prompted, enter y.</p> <p>f. Return to the admin level by entering the following command:</p> <pre>set -privilege admin</pre> <p>If necessary, you can force the relocation using one of the following methods:</p> <ul style="list-style-type: none"> • Overriding veto checks: <pre>storage aggregate relocation start -override -vetoes -ndo-controller-upgrade</pre> • Overriding destination checks: <pre>storage aggregate relocation start -override -destination-checks -ndocontroller-upgrade</pre> <p>For more information about storage aggregate relocation commands refer to References to link to <i>Disk and aggregate management with the CLI</i> and the <i>ONTAP 9 Commands: Manual Page Reference</i>.</p>

5. Verify that all node2 non-root aggregates are online and their state on node4:

```
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:

```
volume show -node node4 -state offline
```

8. If any volumes are offline on node4, bring them online:

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

9. Send a post-upgrade AutoSupport message to NetApp for node4:

```
system node autosupport invoke -node node4 -type all -message "node2  
successfully upgraded from platform_old to platform_new"
```

Stage 6. Complete the upgrade

Manage authentication using KMIP servers

With ONTAP 9.5 and later, you can use Key Management Interoperability Protocol (KMIP) servers to manage authentication keys.

Steps

1. Add a new controller:

```
security key-manager setup -node new_controller_name
```

2. Add the key manager:

```
security key-manager -add key_management_server_ip_address
```

3. Verify that the key management servers are configured and available to all nodes in the cluster:

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

4. Restore the authentication keys from all linked key management servers to the new node:

```
security key-manager restore -node new_controller_name
```

Confirm that the new controllers are set up correctly

To confirm correct setup, you enable the HA pair. You 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 confirm 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:

```
storage failover show
```

The following example shows the output of the command when storage failover is enabled:

```
cluster::> storage failover show
```

Node	Partner	Takeover Possible	State Description
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: <code>cluster ha modify -configured true</code>
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:

```
storage aggregate show -owner-name <node3>
storage aggregate show -owner-name <node4>
```

8. Determine whether any volumes are offline:

```
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 [Step 19 \(d\)](#) in *Prepare the nodes for upgrade*, and bring online any of the offline volumes, as required, by entering the following command, once for each volume:

```
volume online -vserver <vserver_name> -volume <volume_name>
```

10. Install new licenses for the new nodes by entering the following command 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 multiple licenses at once, each license key separated by a comma.

11. If self-encrypting drives are being used in the configuration and you have set the `kmip.init.maxwait` variable to `off` (for example, in [Step 16](#) of *Install and boot node3*), you must unset the variable:

```
set diag; systemshell -node node_name -command sudo kenv -u -p  
kmip.init.maxwait
```

12. 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.

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

```
system license show
```

You can compare the output with the output that you captured in [Step 29](#) of *Prepare the nodes for upgrade*.

14. Configure the SPs by performing the following command on both nodes:

```
system service-processor network modify -node <node_name>
```

Go to [References](#) to link to the *System Administration Reference* for information about the SPs and the *ONTAP 9 Commands: Manual Page Reference* for detailed information about the `system service-`

processor network modify command.

15. If you want to set up a switchless cluster on the new nodes, go to [References](#) to link to the *Network Support Site* and follow the instructions in *Transitioning to a two-node switchless cluster*.

After you finish

If Storage Encryption is enabled on node3 and node4, complete the steps in [Set up Storage Encryption on the new controller module](#). Otherwise, complete the steps in [Decommission the old system](#).

Set 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. You must 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 the previous step 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 the 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
```

Set up NetApp Volume or Aggregate Encryption on the new controller module

If the replaced controller or the high-availability (HA) partner of the new controller uses NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE), you must configure the new controller module for NVE or NAE.

About this task

This procedure includes steps that are performed on the new controller module. You must 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:

For this ONTAP version...	Use this command...
ONTAP 9.6 or 9.7	<code>security key-manager key query -node node</code>
ONTAP 9.5 or earlier	<code>security key-manager key show</code>

2. Add the key management servers listed in the previous step to the key management server list in the new controller:

- a. Add the key management server by using the following command:

```
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 by using the following command:

```
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 by using the following command:

```
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.

- Restore authentication for external key manager:

```
security key-manager external restore
```

This command needs the Onboard Key Manager (OKM) passphrase

For more information, see the Knowledge Base article [How to restore external key manager server configuration from the ONTAP boot menu](#).

- Restore authentication for the OKM:

For this ONTAP version...	Use this command...
All other ONTAP versions	<code>security key-manager onboard sync</code>
ONTAP 9.5	<code>security key-manager setup -node <i>node_name</i></code>

After you finish

Check if any volumes were taken offline because authentication keys were not available or External Key Management servers could not be reached. Bring those volumes back online by using the `volume online` command.

Decommission 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. Refer to [References](#) to link to the *NetApp Support Site* and log in.
2. Select **Products > My Products** from the menu.
3. On the **View Installed Systems** page, choose which **Selection Criteria** you want to use to display information about your system.

You can choose one of the following to locate your system:

- Serial Number (located on the back of the unit)
- Serial Numbers for My Location

4. Select **Go!**

A table displays cluster information, including the serial numbers.

5. Locate the cluster in the table and select **Decommission this system** from the Product Tool Set drop-down menu.

Resume SnapMirror operations

You can resume SnapMirror transfers that were quiesced before upgrade and resume the SnapMirror relationships. The updates are on schedule after the upgrade is completed.

Steps

1. Verify the SnapMirror status on the destination:

```
snapmirror show
```

2. Resume the SnapMirror relationship:

```
snapmirror resume -destination-vserver vsilver_name
```

Troubleshoot

Aggregate relocation failures

Aggregate relocation (ARL) might fail at different points during the upgrade.

Check for aggregate relocation failure

During the procedure, ARL might fail in Stage 2, Stage 3, or Stage 5.

Steps

1. Enter the following command and examine the output:

```
storage aggregate relocation show
```

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, refer to [References](#) to link to the *ONTAP 9 Commands: Manual Page Reference*.

Aggregates originally on node1 are owned by node4 after completion of the upgrade

At the end of the upgrade procedure, node3 must 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 correctly, 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 [Prepare the nodes for 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 [Prepare the nodes for upgrade](#) and note any aggregates that were not correctly relocated.
3. Relocate the aggregates left behind on node4:

```
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 looks 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 *Relocate non-root aggregates from node1 to node2* again.
2. Continue with the node-pair upgrade procedure.

Node1 crashes after Stage 2 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 fails 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 gives back all of node2's aggregates.

Steps

1. Complete [Step 1](#) in the section *Relocate 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 the rest of the node pair upgrade procedure.

Reboots, panics, or power cycles during Stage 3

Failures 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 5](#) in the section *Relocate non-root 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 does 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 5](#) again in the section *Relocate non-root 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 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 *Relocate 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 Relocate 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 during 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	
		Possible	State Description
node1	node2	false	Unknown
node2	node1	false	Node owns partner aggregates as part of the non-disruptive head upgrade procedure. Takeover is not possible: Storage failover is disabled.

The following example shows the output of the storage failover show command after a reboot or panic:

```
cluster::> storage failover show
```

Node	Partner	Takeover	
		Possible	State Description
node1	node2	-	Unknown
node2	node1	false	Waiting for node1, Partial giveback, Takeover is not possible: Storage failover is disabled

Although the output says that a node is in partial giveback and that storage failover is disabled, you can disregard this message.

Steps

No action is required; continue with the node-pair upgrade procedure.

LIF migration failure

After you migrate LIFs, they might not come online after migration in Stage 2, Stage 3, or Stage 5.

Steps

1. Verify that the 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.
2. Check the physical connectivity of the network cable if the physical state of the port is "down".

LIFs are on invalid ports after upgrade

After the upgrade is completed, the FC logical interfaces (LIFs) might be left on incorrect ports if you have a MetroCluster configuration. You can perform a resync operation to reassign the LIFs to the correct ports.

Step

1. Enter the `metrocluster vsync resync` command to reallocate the LIFs to the correct ports.


```
metrocluster vsync resync -vsync vsync_name fcp-mc.headupgrade.test.vsync
```

References

When performing the procedures in this content, you might need to consult reference content or go to reference websites.

Reference content

Content specific to this upgrade are listed in the table below.

Content	Description
Administration overview with the CLI	Describes how to administer ONTAP systems, shows you how to use the CLI interface, how to access the cluster, how to manage nodes, and much more.
Decide whether to use System Manager or the ONTAP CLI for cluster setup	Describes how to set up and configure ONTAP.
Disk and aggregate management with the CLI	Describes how to manage ONTAP physical storage using the CLI. It shows you how to create, expand, and manage aggregates, how to work with Flash Pool aggregates, how to manage disks, and how to manage RAID policies.

Content	Description
HA pair management	Describes how to install and manage high-availability clustered configurations, including storage failover and takeover/giveback.
Logical storage management with the CLI	Describes how to efficiently manage your logical storage resources, using volumes, FlexClone volumes, files, and LUNs, FlexCache volumes, deduplication, compression, qtrees, and quotas.
MetroCluster Management and Disaster Recovery	Describes how to perform MetroCluster switchover and switchback operations, both in planned maintenance operations, or in the event of a disaster.
MetroCluster Upgrade and Expansion	Provides procedures for upgrading controller and storage models in the MetroCluster configuration, transitioning from a MetroCluster FC to a MetroCluster IP configuration, and expanding the MetroCluster configuration by adding additional nodes.
Network Management	Describes how to configure and manage physical and virtual network ports (VLANs and interface groups), LIFs, routing, and host-resolution services in clusters; optimize network traffic by load balancing; and monitor the cluster by using SNMP.
ONTAP 9.0 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.0 commands.
ONTAP 9.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.1 commands.
ONTAP 9.2 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.2 commands.
ONTAP 9.3 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.3 commands.
ONTAP 9.4 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.4 commands.
ONTAP 9.5 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.5 commands.
ONTAP 9.6 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.6 commands.
ONTAP 9.7 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.7 commands.
ONTAP 9.8 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.8 commands.
ONTAP 9.9.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.9.1 commands.
ONTAP 9.10.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.10.1 commands.
SAN management with the CLI	Describes how to configure and manage LUNs, igroups, and targets using the iSCSI and FC protocols, and namespaces and subsystems using the NVMe/FC protocol.

Content	Description
SAN configuration reference	Contains information about FC and iSCSI topologies and wiring schemes.
Upgrade by moving volumes or storage	Describes how to quickly upgrade controller hardware in a cluster by moving storage or volumes. Also describes how to convert a supported model to a disk shelf.
Upgrade ONTAP	Contains instructions for downloading and upgrading ONTAP.
Use "system controller replace" commands to upgrade controller hardware introduced in ONTAP 9.15.1 and later	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers introduced in ONTAP 9.15.1 and later by using "system controller replace" commands.
Use "system controller replace" commands to upgrade controller models in the same chassis	Describes the aggregate relocation procedures needed to non-disruptively upgrade a system, keeping the old system chassis and disks.
Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers running ONTAP 9.8 by using "system controller replace" commands.
Use aggregate relocation to manually upgrade controller hardware running ONTAP 9.8 or later	Describes the aggregate relocation procedures needed to perform manual non-disruptive controller upgrades running ONTAP 9.8 or later.
Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.5 to ONTAP 9.7	Describes the aggregate relocation procedures needed to non-disruptively upgrade controllers running ONTAP 9.5 to ONTAP 9.7 by using "system controller replace" commands.
Use aggregate relocation to manually upgrade controller hardware running ONTAP 9.7 or earlier	Describes the aggregate relocation procedures needed to perform manual non-disruptive controller upgrades running ONTAP 9.7 or earlier.

Reference sites

The [NetApp Support Site](#) also contains documentation about network interface cards (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.

Access [ONTAP 9 documentation](#).

Access the [Active IQ Config Advisor](#) tool.

Copyright information

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

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

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

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

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

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

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

Trademark information

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