



Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.5 to 9.7

Upgrade controllers

NetApp
February 19, 2026

Table of Contents

- Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.5 to 9.7 1
 - Learn about this ARL upgrade procedure 1
 - Automate the controller upgrade process 1
 - Decide whether to use this aggregate relocation procedure 2
 - Supported system upgrade combinations 2
 - Choose a different hardware upgrade procedure 3
 - Required tools and documentation 3
 - Guidelines for upgrading controllers with ARL 4
 - Supported upgrades for ARL 4
 - Two-node switchless clusters 4
 - Upgrades not supported for ARL 4
 - MetroCluster FC configuration 5
 - Troubleshoot 5
 - Verify the health of the MetroCluster configuration 5
 - Check for MetroCluster configuration errors 6
 - Verify switchover, healing, and switchback 6
 - Learn about the ARL upgrade sequence 7
 - Upgrade the node pair 7
 - ARL upgrade sequence overview 7
 - Stage 1. Prepare for upgrade 8
 - Prepare the nodes for upgrade 8
 - Manage storage encryption using the Onboard Key Manager 13
 - Stage 2. Relocate and retire node1 14
 - Relocate non-root aggregates and NAS data LIFs owned by node1 to node2 14
 - Relocate failed or vetoed aggregates 16
 - Retire node1 17
 - Prepare for netboot 17
 - Stage 3. Install and boot node3 18
 - Install and boot node3 18
 - Set the FC or UTA/UTA2 configuration on node3 23
 - Reassign node1 disks to node3 27
 - Map ports from node1 to node3 32
 - Join the quorum when a node has a different set of network ports 36
 - Verify the node3 installation 38
 - Move non-root aggregates and NAS data LIFs owned by node1 from node2 to node3 39
 - Stage 4. Relocate and retire node2 42
 - Relocate non-root aggregates and NAS data LIFs from node2 to node3 42
 - Retire node2 45
 - Stage 5. Install and boot node4 45
 - Install and boot node4 45
 - Set the FC or UTA/UTA2 configuration on node4 50
 - Reassign node2 disks to node4 54
 - Map ports from node2 to node4 59

Join the quorum when a node has a different set of network ports	63
Verify the node4 installation	64
Move non-root aggregates and NAS data LIFs owned by node2 from node3 to node4	66
Stage 6. Complete the upgrade	69
Manage authentication using KMIP servers.	69
Confirm that the new controllers are set up correctly.	70
Set up Storage Encryption on the new controller module	72
Set up NetApp Volume or Aggregate Encryption on the new controller module	73
Decommission the old system	76
Resume SnapMirror operations	76
Troubleshoot	76
Aggregate relocation failures	76
Reboots, panics, or power cycles	78
Issues that can arise in multiple stages of the procedure	81
LIF migration failure	82
References	83
Reference content	83
Reference sites	85

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

Old controller	Replacement controller
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
```

```
dpggqa-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 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 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.

Stage	Steps
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, 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.
Stage 4. Relocate and retire node2	<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.
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, 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 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 NetApp Volume Encryption. 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 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.

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.
- 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:

- ° `vserver 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.

- 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:

- a. 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 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: <pre> storage aggregate relocation start -override -vetoes * -ndocontroller-upgrade true </pre>
Overriding destination checks	Enter the following: <pre> storage aggregate relocation start -overridedestination-checks * -ndo -controllerupgrade true </pre>

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

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	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> <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.  Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.

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	<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 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
      Current  Current  Pending  Pending  Admin
Adapter Mode    Type    Mode    Type    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 .

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.

10. Exit maintenance mode:

```
halt
```

11. 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

1. 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 -ip-space 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 vservers -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 -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
```

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 -ip-space 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 ip-space 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 vsserver -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 -ip-space 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|fcp -home-node node3
```

The system returns output similar to the following example:

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

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

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

i. Set the LIF status to down:

```
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 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](#).

- 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	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> <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. NOTE: Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.

- 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. 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 -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
```

- 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 -ip-space Cluster -broadcast
-domain Cluster -ports node4:e0d
```

- h. Verify that node4 has rejoined quorum:

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

5. Adjust the broadcast domains hosting your cluster LIFs and node-management/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 -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
```

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 -ip-space 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 vserver-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 node_name</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 lif_name -status
-admin down

```

ii. Remove the LIF from the port set:

```

portset remove -vserver vs0 -portset portset_name -port-name
port_name

```

iii. Enter one of the following commands:

- Move a single LIF by entering the following command:

```

network interface modify -vserver vs0 -lif lif_name -home
-port new_home_port

```

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

```

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

```

- Add the LIFs back to the port set:

```

portset add -vserver vs0 -portset portset_name -port-name
port_name

```



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 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 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
-----  -
node1     node2      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
-----  -
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.

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.