



Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later

Upgrade controllers

NetApp
August 14, 2025

Table of Contents

| | |
|--|----|
| Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later | 1 |
| Learn about this ARL upgrade procedure | 1 |
| Automate the controller upgrade process | 2 |
| Decide whether to use this aggregate relocation procedure | 2 |
| Supported system upgrade combinations | 3 |
| Choose a different hardware upgrade procedure | 4 |
| Required tools and documentation | 4 |
| Guidelines for upgrading controllers with ARL | 5 |
| Supported upgrades for ARL | 5 |
| Two-node switchless clusters | 5 |
| Upgrades not supported for ARL | 5 |
| MetroCluster FC configuration | 6 |
| Troubleshoot | 6 |
| Verify the health of the MetroCluster configuration | 6 |
| Check for MetroCluster configuration errors | 7 |
| Verify switchover, healing, and switchback | 7 |
| Learn about the ARL upgrade sequence | 8 |
| Upgrade the node pair | 8 |
| ARL upgrade sequence overview | 8 |
| Stage 1. Prepare for upgrade | 9 |
| Prepare the nodes for upgrade | 10 |
| Manage storage encryption using the Onboard Key Manager | 14 |
| Stage 2. Relocate and retire node1 | 15 |
| Relocate non-root aggregates and NAS data LIFs owned by node1 to node2 | 15 |
| 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 | 25 |
| Reassign node1 disks to node3 | 30 |
| Verify the node3 installation | 37 |
| Restore key-manager configuration on node3 | 44 |
| Move non-root aggregates and NAS data LIFs owned by node1 from node2 to node3 | 45 |
| Stage 4. Relocate and retire node2 | 47 |
| Relocate non-root aggregates and NAS data LIFs from node2 to node3 | 47 |
| Retire node2 | 48 |
| Stage 5. Install and boot node4 | 48 |
| Install and boot node4 | 48 |
| Set the FC or UTA/UTA2 configuration on node4 | 55 |
| Reassign node2 disks to node4 | 60 |
| Verify the node4 installation | 66 |
| Restore key-manager configuration on node4 | 74 |

| | |
|---|----|
| Move non-root aggregates and NAS data LIFs owned by node2 from node3 to node4 | 75 |
| Stage 6. Complete the upgrade | 77 |
| Manage authentication using KMIP servers | 77 |
| Confirm that the new controllers are set up correctly | 77 |
| Set up Storage Encryption on the new controller module | 80 |
| Set up NetApp Volume or Aggregate Encryption on the new controller module | 81 |
| Decommission the old system | 83 |
| Resume SnapMirror operations | 83 |
| Troubleshoot | 83 |
| Aggregate relocation failures | 83 |
| Reboots, panics, or power cycles | 85 |
| Issues that can arise in multiple stages of the procedure | 88 |
| LIF migration failure | 89 |
| References | 90 |
| Reference content | 90 |
| Reference sites | 92 |

Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later

Learn about this ARL upgrade procedure

There are several aggregate relocation (ARL) methods for upgrading controller hardware. This procedure describes how to upgrade the controller hardware using aggregate relocation (ARL) with "system controller replace commands" on systems running ONTAP 9.8 or later.

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

Terminology used in this information

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

Throughout this information, the term "systems with FlexArray Virtualization Software" refers to systems that belong to these new platforms. The term "V-Series system" refers to the separate hardware systems that can attach to storage arrays.

Important information:

- This procedure is complex and assumes that you have advanced ONTAP administration skills. You also must read and understand the [guidelines for upgrading controllers with ARL](#) and the [ARL upgrade sequence](#) before beginning the upgrade.
- This procedure assumes that the replacement controller hardware is new and has not been used. The steps required to prepare used controllers with the `wipeconfig` command are not included in this procedure. You must contact technical support if the replacement controller hardware was previously used, especially if the controllers were running Data ONTAP in 7- Mode.
- You can use ARL to perform a non-disruptive simplified controller upgrade to a new controller running a later ONTAP version than the version running on the cluster you are upgrading. The ONTAP version combinations for old and new controllers are determined by the ONTAP software release NDU cadence model. For example, if you have a controller running ONTAP 9.8, and that is the last supported version for that controller, you can upgrade to a new controller running an ONTAP version later than ONTAP 9.8.

This upgrade procedure primarily applies to upgrade scenarios where the controller model you are replacing does not support later ONTAP versions and the new controller does not support earlier ONTAP versions.

- You can use this procedure to upgrade the controller hardware in clusters with more than two nodes; however, you need to perform the procedure separately for each high-availability (HA) pair in the cluster.

- This procedure applies to FAS systems, V-Series systems, AFF systems, and systems with FlexArray Virtualization Software. FAS systems released after ONTAP 9.5 can attach to storage arrays if the required license is installed. For more information about the storage array and V-Series models, refer to [References](#) to link to the *Hardware Universe* and go to the V-Series Support Matrix.
- This procedure applies to systems running 4-node NetApp MetroCluster configuration or higher. Since MetroCluster configuration sites can be at two physically different locations, the automated controller upgrade must be carried out individually at each MetroCluster site for an HA pair.
- For non-MetroCluster systems, such as HA clusters, the ARL upgrade is the only supported procedure.
- If you are upgrading from an AFF A320 system, you can use volume moves to upgrade controller hardware or contact technical support. Refer to [References](#) to link to *Upgrade by moving volumes or storage*.

Automate the controller upgrade process

During a controller upgrade, the controller is replaced with another controller running a newer or more powerful platform. Earlier versions of this content contained instructions for a nondisruptive controller update process that was comprised of entirely manual steps. This content provides the steps for the new automated procedure which utilizes automatic network port reachability checks to further simplify the controller upgrade experience.

The manual process was lengthy and complex but in this simplified procedure you can implement a controller update using aggregate relocation, enabling more efficient nondisruptive upgrades for HA pairs. There are significantly fewer manual steps, especially around validation, collection of information, and post checks.

Decide whether to use this aggregate relocation procedure

There are several aggregate relocation (ARL) methods for upgrading controller hardware. This procedure describes how to upgrade the controller hardware using aggregate relocation (ARL) with "system controller replace commands" on systems running ONTAP 9.8 or later. You should only use this complex procedure if you're an experienced ONTAP administrator.

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

- You're running ONTAP 9.8 or later.
- You don't want to add the new controllers as a new HA pair to the cluster and migrate the data using volume moves.
- You're experienced in administering ONTAP and are comfortable with the risks of working in the diagnostic privilege mode.
- If you're upgrading a MetroCluster configuration, it's a four-node or higher FC configuration, and all nodes are running ONTAP 9.8 or later.

For upgrading MetroCluster IP configurations, refer to [References](#) to link to *MetroCluster Upgrade and Expansion*.



- If you're upgrading a system by swapping controller modules in the same chassis, such as AFF A800 or AFF C800, NetApp strongly recommends using the upgrade procedure that [upgrades controller models using ARL, keeping the existing system chassis and disks](#). This ARL procedure includes the steps that ensure the internal disks remain secure in the chassis when you remove and install the controllers during the upgrade procedure.

[Learn about the supported system upgrade combinations using ARL, keeping the existing system chassis and disks.](#)

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

Supported system upgrade combinations

The following table shows the supported systems matrix for performing a controller upgrade using this ARL procedure.

| Old controller | Replacement controller |
|--|--|
| FAS8020 ³ , FAS8040 ³ , FAS8060, FAS8080 | FAS8200, FAS8300, FAS8700, FAS9000 |
| FAS8060 ⁴ , FAS8080 ⁴ | FAS9500 |
| AFF8020 ³ , AFF8040 ³ , AFF8060, AFF8080 | AFF A300, AFF A400, AFF A700, AFF A800 ¹ |
| AFF8060 ⁴ , AFF8080 ⁴ | AFF A900 |
| FAS8200 | FAS8300 ² , FAS8700, FAS9000, FAS9500 |
| FAS8300, FAS8700, FAS9000 | FAS9500 |
| AFF A300 | AFF A400 ² , AFF A700, AFF A800 ¹ , AFF A900 |
| AFF A320 ⁴ | AFF A400 |
| AFF A400, AFF A700 | AFF A900 |



If your controller upgrade model combination isn't in the above table, contact technical support.

¹For the additional steps needed for AFF A800 systems, go to the step that references A800 in section [Reassign node1 disks to node3, Step 9](#), or [Reassign node2 disks to node4, Step 9](#).

²If you're upgrading from an AFF A300 to an AFF A400 or an FAS8200 to an FAS8300 system in a two-node switchless cluster configuration, you must pick temporary cluster ports for the controller upgrade. The AFF A400 and FAS8300 systems come in two configurations, as an Ethernet bundle where the mezzanine card ports are Ethernet type and as an FC bundle where the mezzanine ports are FC type.

- For an AFF A400 or an FAS8300 with an Ethernet type configuration, you can use any of the two mezzanine ports as temporary cluster ports.
- For an AFF A400 or an FAS8300 with an FC type configuration, you must add a four-port 10GbE network interface card (part number X1147A) to provide temporary cluster ports.
- After you complete a controller upgrade by using temporary cluster ports, you can nondisruptively migrate cluster LIFs to e3a and e3b, 100GbE ports on an AFF A400 system, and e0c and e0d, 100GbE ports on an FAS8300 system.

³For FAS8020, FAS8040, AFF8020, and AFF8040 system upgrades to the target replacement controllers listed

in the table above, the replacement controllers must be running same ONTAP version as the old controller. Note that FAS8020, FAS8040, AFF8020, and AFF8040 systems do not support ONTAP versions later than ONTAP 9.8.

⁴The following table shows the minimum and later supported ONTAP versions for these controller upgrade combinations.

| Old controller | | Replacement controller | |
|----------------|-------------------------|------------------------|--------------------|
| System | ONTAP version | System | ONTAP version |
| AFF A320 | 9.9.1 or later | AFF A400 | 9.9.1 or later |
| AFF8060 | 9.8P13 or later patches | AFF A900 | 9.10.1 to 9.12.1 |
| AFF8080 | 9.8P10 or later patches | AFF A900 | 9.10.1 to 9.12.1 |
| FAS8060 | 9.8P13 or later patches | FAS9500 | 9.10.1P3 to 9.12.1 |
| FAS8080 | 9.8P12 or later patches | FAS9500 | 9.10.1P3 to 9.12.1 |

For the upgrade combinations shown in the preceding table:



- It isn't required to use the same ONTAP version on the existing and replacement controllers. The ONTAP software upgrade is performed with the controller upgrade.
- When upgrading, you must install a replacement controller with a supported ONTAP version and patch level.
- It isn't possible to cancel or back out of a controller upgrade after you start the procedure and upgrade the first node.

Choose a different hardware upgrade procedure

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

Related information

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

Required tools and documentation

You must have specific tools to install the new hardware, and you need to reference other documents during the upgrade process.

You need the following tools to perform the upgrade:

- Grounding strap
- #2 Phillips screwdriver

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

Guidelines for upgrading controllers with ARL

To understand whether you can use ARL to upgrade a pair of controllers running ONTAP 9.8 or later depends on the platform and the configuration of both the original and replacement controllers.

Supported upgrades for ARL

When you upgrade a pair of nodes using this ARL procedure for ONTAP 9.8 or later, you must verify that ARL can be performed on the original and replacement controllers.

You must check the size of all defined aggregates and number of disks supported by the original system. You must then compare the aggregate sizes and number of disks supported to the aggregate size and number of disks supported by the new system. Refer to [References](#) to link to the *Hardware Universe* where this information is available. The aggregate size and the number of disks supported by the new system must be equal to or greater than the aggregate size and number of disks supported by the original system.

You must validate in the cluster mixing rules whether new nodes can become part of the cluster with the existing nodes when the original controller is replaced. For more information about cluster mixing rules, refer to [References](#) to link to the *Hardware Universe*.



If you are upgrading a system that supports internal drives (for example, an FAS2700 or AFF A250) but does NOT have internal drives, refer to [References](#) and use the procedure in the *Aggregate Relocation to Manually Upgrade Controller Hardware* content that is correct for your version of ONTAP.

If you have a system with more than two cluster ports per node, such as an FAS8080 or an AFF8080 system, before you start the upgrade, you must migrate and re-home the cluster LIFs to two cluster ports per node. If you perform the controller upgrade with more than two cluster ports per node, cluster LIFs might be missing on the new controller after the upgrade.

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

Two-node switchless clusters

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

Upgrades not supported for ARL

You cannot perform the following upgrades:

- To replacement controllers that do not support the disk shelves connected to the original controllers

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

- To entry level controllers with internal drives, for example: an FAS 2500.

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

MetroCluster FC configuration

In a MetroCluster FC configuration, you must replace the disaster recovery/failover site nodes as soon as possible. Mismatch in controller models within a MetroCluster is not supported because controller model mismatch can cause disaster recovery mirroring to go offline. Use the `-skip-metrocluster-check true` command to bypass MetroCluster checks when you are replacing nodes at the second site.

Troubleshoot

You might encounter a failure while upgrading the node pair. The node might crash, aggregates might not relocate, or LIFs might not migrate. The cause of the failure and its solution depend on when the failure occurred during the upgrade procedure.

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

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

Verify the health of the MetroCluster configuration

Before starting an upgrade on a Fabric MetroCluster configuration, you must check the health of the MetroCluster configuration to verify correct operation.

Steps

1. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
metrocluster_siteA::*> metrocluster check run
```

The operation runs in the background.

2. After the `metrocluster check run` operation completes, view the results:

```
metrocluster check show
```

After approximately five minutes, the following results are displayed:

```
metrocluster_siteA:*> metrocluster check show
Last Checked On: 4/7/2019 21:15:05
Component          Result
-----
nodes              ok
lifs               ok
config-replication ok
aggregates         warning
clusters           ok
connections        not-applicable
volumes            ok
7 entries were displayed.
```

3. Check the status of the running MetroCluster check operation:

```
metrocluster operation history show -job-id 38
```

4. Verify that there are no health alerts:

```
system health alert show
```

Check for MetroCluster configuration errors

You can use the Active IQ Config Advisor tool available from the NetApp Support Site to check for common configuration errors.

If you do not have a MetroCluster configuration, you can skip this section.

About this task

Active IQ Config Advisor is a configuration validation and health check tool. You can deploy it at both secure sites and non-secure sites for data collection and system analysis.



Support for Config Advisor is limited, and available only online.

1. Download the [Active IQ Config Advisor](#) tool.
2. Run Active IQ Config Advisor, reviewing the output and following its recommendations to address any issues.

Verify switchover, healing, and switchback

You should verify the switchover, healing, and switchback operations of the MetroCluster configuration.

Refer to [References](#) to link to the *MetroCluster Management and Disaster Recovery* content and use the procedures mentioned for negotiated switchover, healing, and switchback.

Learn about the ARL upgrade sequence

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

Upgrade the node pair

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

ARL upgrade sequence overview

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

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

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

| Stage | Description |
|--|---|
| Stage 1. Prepare for upgrade | <p>During Stage 1, you run prechecks and, if required, correct aggregate ownership. You must record certain information if you are managing storage encryption by using the OKM and you can choose to quiesce the SnapMirror relationships.</p> <p>Aggregate ownership at the end of Stage 1:</p> <ul style="list-style-type: none">• Node1 is the home owner and current owner of the node1 aggregates.• Node2 is the home owner and current owner of the node2 aggregates. |
| Stage 2. Relocate and retire node1 | <p>During Stage 2, you relocate node1 non-root aggregates and NAS data LIFs to node2. This process is largely automated; the operation pauses to enable you to check its status. You must manually resume the operation. If required, you relocate failed or vetoed aggregates. You record node1 information for use later in the procedure before retiring node1. You can also prepare to netboot node3 and node4 later in the procedure.</p> <p>Aggregate ownership at the end of Stage 2:</p> <ul style="list-style-type: none">• Node2 is the current owner of node1 aggregates.• Node2 is the home owner and current owner of node2 aggregates. |

| Stage | Description |
|------------------------------------|---|
| Stage 3. Install and boot node3 | <p>During Stage 3, you install and boot node3, check that the cluster and node-management ports from node1 come online on node3, reassign the node1 disks to node3, and verify the node3 installation. If you are using NetApp Volume Encryption (NVE), you restore key-manager configuration. If required, you set the FC or UTA/UTA2 configuration on node3. You also relocate the node1 NAS data LIFs and non-root aggregates from node2 to node3 and verify that the SAN LIFs exist on node3.</p> <p>Aggregate ownership at the end of Stage 3:</p> <ul style="list-style-type: none"> • Node3 is the home owner and current owner of node1 aggregates. • Node2 is the home owner and current owner of node2 aggregates. |
| Stage 4. Relocate and retire node2 | <p>During Stage 4, you relocate non-root aggregates and NAS data LIFs from node2 to node3. You also record node2 information for use later in the procedure before retiring node2.</p> <p>Aggregate ownership at the end of Stage 4:</p> <ul style="list-style-type: none"> • Node3 is the home owner and current owner of aggregates that originally belonged to node1. • Node2 is the home owner of node2 aggregates. • Node3 is the current owner of node2 aggregates. |
| Stage 5. Install and boot node4 | <p>During Stage 5, you install and boot node4, check that the cluster and node-management ports from node2 come online on node4, reassign the node2 disks to node4, and verify the node4 installation. If you are using NVE, you restore key-manager configuration. If required, you set the FC or UTA/UTA2 configuration on node4. You also relocate node2 NAS data LIFs and non-root aggregates from node3 to node4 and verify that the SAN LIFs exist on node4.</p> <p>Aggregate ownership at the end of Stage 5:</p> <ul style="list-style-type: none"> • Node3 is the home owner and current owner of the aggregates that originally belonged to node1. • Node4 is the home owner and current owner of aggregates that originally belonged to node2. |
| Stage 6. Complete the upgrade | <p>During Stage 6, you confirm that the new nodes are set up correctly and, if the new nodes are encryption-enabled, you configure and set up Storage Encryption or NVE. You should also decommission the old nodes and resume the SnapMirror operations.</p> |

Stage 1. Prepare for upgrade

Prepare the nodes for upgrade

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

Steps

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

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



- Beginning with ONTAP 9.10.1, the automated negotiated switchover (NSO) based upgrade procedure is the default for a four-node MetroCluster FC configuration. If you are upgrading a four-node MetroCluster FC configuration, when you issue the `system controller replace start` command, you must prevent the NSO based procedure initiating by setting the `-nso` parameter to `false`:

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

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

```
set -privilege advanced
```

You will see the following output:

Warning:

1. Current ONTAP version is 9.x

Before starting controller replacement operation, ensure that the new controllers are running the version 9.x

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

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

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

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

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

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

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

| Precheck | Description |
|-----------------------------------|---|
| Cluster Health Check | Checks all the nodes in the cluster to confirm they are healthy. |
| MCC Cluster Check | Checks if the system is a MetroCluster configuration. The operation automatically detects if it is a MetroCluster configuration or not and performs the specific prechecks and verification checks. Only 4-node MetroCluster FC configuration is supported. In the case of 2-node MetroCluster configuration and 4-node MetroCluster IP configuration, the check fails. If the MetroCluster configuration is in switched over state, the check fails. |
| Aggregate Relocation Status Check | Checks whether an aggregate relocation is already in progress. If another aggregate relocation is in progress, the check fails. |
| Model Name Check | Checks whether the controller models are supported for this procedure. If the models are not supported, the task fails. |
| Cluster Quorum Check | Checks that the nodes being replaced are in quorum. If the nodes are not in quorum, the task fails. |
| Image Version Check | Checks that the nodes being replaced run the same version of ONTAP. If the ONTAP image versions are different, the task fails. The new nodes must have the same version of ONTAP 9.x installed on them that is installed on the original nodes. If the new nodes have a different version of ONTAP installed, you need to netboot the new controllers after you install them. For instructions on how to upgrade ONTAP, refer to References to link to <i>Upgrade ONTAP</i> . |
| HA Status Check | Checks if both the nodes being replaced are in a high-availability (HA) pair configuration. If storage failover is not enabled for the controllers, the task fails. |
| Aggregate Status Check | If the nodes being replaced own aggregates for which they are not the home owner, the task fails. The nodes should not own any non-local aggregates. |
| Disk Status Check | If any nodes being replaced have missing or failed disks, the task fails. If any disks are missing, refer to References to link to <i>Disk and aggregate management with the CLI</i> , <i>Logical storage management with the CLI</i> , and <i>HA pair management</i> to configure storage for the HA pair. |

| Precheck | Description |
|--------------------------------|---|
| Data LIF Status Check | Checks if any of the nodes being replaced have non- local data LIFs. The nodes should not contain any data LIFs for which they are not the home owner. If one of the nodes contains non-local data LIFs, the task fails. |
| Cluster LIF Status | Checks whether the cluster LIFs are up for both nodes. If the cluster LIFs are down, the task fails. |
| ASUP Status Check | If ASUP notifications are not configured, the task fails. You must enable ASUP before beginning the controller replacement procedure. |
| CPU Utilization Check | Checks if the CPU utilization is more than 50% for any of the nodes being replaced. If the CPU usage is more than 50% for a considerable period of time, the task fails. |
| Aggregate Reconstruction Check | Checks if reconstruction is occurring on any data aggregates. If aggregate reconstruction is in progress, the task fails. |
| Node Affinity Job Check | Checks if any node affinity jobs are running. If node affinity jobs are running, the check fails. |

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



If you have a system with more than two cluster ports per node, such as an FAS8080 or an AFF8080 system, before you start the upgrade, you must migrate and re-home the cluster LIFs to two cluster ports per node. If you perform the controller upgrade with more than two cluster ports per node, cluster LIFs might be missing on the new controller after the upgrade.

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

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

```

° vservers services name-service dns show
° network interface show -curr-node local -role cluster,intercluster,node-
  mgmt,cluster-mgmt,data
° network port show -node local -type physical
° service-processor show -node local -instance
° network fcp adapter show -node local
° network port ifgrp show -node local
° system node show -instance -node local
° run -node local sysconfig
° storage aggregate show -node local

```

- `volume show -node local`
- `storage array config show -switch switch_name`
- `system license show -owner local`
- `storage encryption disk show`
- `security key-manager onboard show-backup`
- `security key-manager external show`
- `security key-manager external show-status`
- `network port reachability show -detail -node local`



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

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

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



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

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

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

Correct aggregate ownership if an ARL precheck fails

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

Steps

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

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

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

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

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


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

4 entries were displayed.
```

After you finish

You must restart the controller replacement process:

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

License

Some features require licenses, which are issued as *packages* that include one or more features. Each node in the cluster must have its own key for each feature to be used in the cluster.

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

Refer to [References](#) to link to the *NetApp Support Site* where you can obtain new 28-character license keys for ONTAP. The keys are available in the *My Support* section under *Software licenses*. If the site does not have the license keys you need, you can contact your NetApp sales representative.

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

Manage storage encryption using the Onboard Key Manager

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

Steps

1. Record the cluster-wide passphrase.

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

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

Quiesce the SnapMirror relationships (optional)

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

Steps

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

```
snapmirror show
```



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

```
snapmirror abort -destination-vserver vservers_name
```

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

2. Quiesce all relationships between the cluster:

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

Stage 2. Relocate and retire node1

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

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

Before you begin

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

About this task

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



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

Steps

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

```
system controller replace resume
```

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

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

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

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

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

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

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

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

| Aggregate | Size | Available | Used% | State | #Vols | Nodes | RAID Status |
|----------------|---------|-----------|-------|--------|-------|-------|-------------|
| aggr_1 | 744.9GB | 744.8GB | 0% | online | 5 | node2 | |
| raid_dp,normal | | | | | | | |
| aggr_2 | 825.0GB | 825.0GB | 0% | online | 1 | node2 | |
| raid_dp,normal | | | | | | | |

2 entries were displayed.

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

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

4. Verify that all the volumes are online on node2 by using the following command on node2 and examining its output:

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

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

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

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

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

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

Relocate failed or vetoed aggregates

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

About this task

The relocation operation will have paused due to the error.

Steps

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

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

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

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

Retire node1

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

Steps

1. Resume the operation:

```
system controller replace resume
```

2. Verify that node1 has been halted:

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

After you finish

You can decommission node1 after the upgrade is completed. See [Decommission the old system](#).

Prepare for netboot

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

Before you begin



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

About this task

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

Steps

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

| For... | Then... |
|----------------------------|---|
| FAS/AFF8000 series systems | <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> <div>  <p>If you are extracting the contents on Windows, use 7-Zip or WinRAR to extract the netboot image.</p> </div> <p>Your directory listing should contain a netboot folder with a kernel file: <code>netboot/kernel</code></p> |
| All other systems | <p>Your directory listing should contain the following file: <code><ontap_version>_image.tgz</code></p> <div>  <p>You do not need to extract the contents of the <code><ontap_version>_image.tgz</code> file.</p> </div> |

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

Stage 3. Install and boot node3

Install and boot node3

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

About this task

The relocation operation is paused at the beginning of this stage. This process is largely automated; the

operation pauses to enable you to check its status. You must manually resume the operation. In addition, you must verify the SAN LIFs have successfully moved to node3.

You need to netboot node3 if it does not have the same version of ONTAP 9 that is installed on node1. After you install node3, boot it from the ONTAP 9 image stored on the web server. You can then download the correct files to the boot media device for subsequent system boots, by following the instructions in [Prepare for netboot](#).



- For an AFF A800 or AFF C800 controller upgrade, you must ensure that all drives in the chassis are firmly seated against the midplane before removing node1. For more information, see [Replace the AFF A800 or AFF C800 controller modules](#).
- If you are upgrading a V-Series system connected to storage arrays or a system with FlexArray Virtualization software that is connected to storage arrays, you need to complete [Step 1](#) through [Step 21](#), then leave this section and follow instructions in the [Configure FC ports on node3](#) and [Check and configure UTA/UTA2 ports on node3](#) sections as needed, entering commands in Maintenance mode. You must then return to this section and resume with [Step 23](#).
- 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 are upgrading to a system with both nodes in the same chassis, install node4 in the chassis as well as node3. If you don't, when you boot node3, the node will behave as if it were in a dual-chassis configuration, and when you boot node4, the interconnect between the nodes will not come up.

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

Cable the following connections, using the *Installation and Setup Instructions* or the *FlexArray Virtualization Installation Requirements and Reference* for the node3 platform, the appropriate disk shelf document, and *HA pair management*.

Refer to [References](#) to link to the *FlexArray Virtualization Installation Requirements and Reference* and *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.



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




Refer to [Prepare for netboot](#).

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



You must use the management port and IP as the netboot connection. Don't use a data LIF IP or a data outage might occur while the upgrade is being performed.

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

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

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. Take one of the following actions:

| If your system... | Description |
|---|---|
| Has disks and no back-end storage | Go to Step 27 |
| Is a V-Series system or a system with FlexArray Virtualization software connected to storage arrays | <div><div>a. Go to section Setting the FC or UTA/UTA2 configuration on node3 and complete the subsections in this section.</div><div>b. Return to this section and complete the remaining steps, beginning with Step 23.</div></div> <div> You must reconfigure FC onboard ports, CNA onboard ports, and CNA cards before you boot ONTAP on the V-Series or system with FlexArray Virtualization software.</div> |

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

If your system has a tape SAN, then you need zoning for the initiators. If required, modify the onboard ports

to initiator by referring to the [Configuring FC ports on node3](#). See your storage array and zoning documentation for further instructions on zoning.

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

See your storage array and zoning documentation for instructions.

25. Modify the worldwide port name (WWPN) values in the host or volume groups associated with array LUNs on the storage array.

Installing a new controller module changes the WWPN values associated with each onboard FC port.

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

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



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

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

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



You cannot mix FIPS drives with other types of drives on the same node or HA pair. You can mix SEDs with non-encrypting drives on the same node or HA pair.

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

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

28. Boot the node into boot menu:

```
boot_ontap menu
```

If you don't have an FC or UTA/UTA2 configuration, perform [Reassign node1 disks to node3, Step 1](#) so that node3 can recognize node1's disks.

29. For a MetroCluster configuration, V-Series systems and systems with FlexArray Virtualization software connected to storage arrays, you must set and configure the FC or UTA/UTA2 ports on node3 to detect the disks attached to the node.

To complete this task, go to section [Set the FC or UTA/UTA2 configuration on 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 does not have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card, and you are upgrading a system with storage disks, you can skip to [Reassign node1 disks to node3](#).
- However, if you have a V-Series system or a system with FlexArray Virtualization software with storage arrays, and node3 does not have onboard FC ports, onboard UTA/UTA ports, or a UTA/UTA2 card, return to the *Install and boot node3* and resume the procedure at [Step 23](#).

Configure FC ports on node3

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

Before you begin

You must have the values of the FC port settings from 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](#).



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

1. Compare the FC settings on node3 with the settings that you captured earlier from node1.
2. Take one of the following actions:

| If the system that you are upgrading... | Then... |
|--|---|
| Has storage disks | <p>In maintenance mode (option 5 at boot menu), modify the FC ports on node3 as needed:</p> <ul style="list-style-type: none"> • To program target ports: <pre>ucadmin modify -m fc -t target adapter</pre> <ul style="list-style-type: none"> • To program initiator ports: <pre>ucadmin modify -m fc -t initiator adapter</pre> <p>-t is the FC4 type: target or initiator.</p> |
| Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays | <p>In maintenance mode (option 5 at boot menu), modify the FC ports on node3 as needed:</p> <pre>ucadmin modify -m fc -t initiator -f adapter_port_name</pre> <p>-t is the FC4 type, target or initiator.</p> <div>  <p>The FC ports must be programmed as initiators.</p> </div> |

- Take one of the following actions:

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

- Exit Maintenance mode:

```
halt
```

- Boot the system from LOADER prompt:

```
boot_ontap menu
```

- After you enter the command, wait until the system stops at the boot environment prompt.
- Select option 5 from the boot menu for maintenance mode.
- Take one of the following actions:

| If the system that you are upgrading... | Then... |
|--|--|
| Has storage disks | <ul style="list-style-type: none"> • If node3 has a UTA/UTA2 card or UTA/UTA2 onboard ports, go to Check and configure UTA/UTA2 ports on node3. • If node3 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports, skip <i>Check and configure UTA/UTA2 ports on node3</i> and go to Reassign node1 disks to node3. |
| Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays | <ul style="list-style-type: none"> • If node3 has a UTA/UTA2 card or UTA/UTA2 onboard ports, go to Check and configure UTA/UTA2 ports on node3. • If node3 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports, skip <i>Check and configure UTA/UTA2 ports on node3</i> and return to <i>Install and boot node3</i> and resume the procedure at Step 23. |

Check and configure UTA/UTA2 ports on node3

If node3 has onboard UTA/UTA2 ports or a UTA/UTA2 card, you must check the configuration of the ports and possibly reconfigure them, depending on how you want to use the upgraded system.

Before you begin

You must have the correct SFP+ modules for the UTA/UTA2 ports.

About this task

If you want to use a Unified Target Adapter (UTA/UTA2) port for FC, you must first verify how the port is configured.



NetApp marketing materials might use the term UTA2 to refer to CNA adapters and ports. However, the CLI uses the term CNA.

You can use the `ucadmin show` command to verify the current port configuration:

```
*> ucadmin show
```

| Adapter | Current Mode | Current Type | Pending Mode | Pending Type | Admin Status |
|---------|--------------|--------------|--------------|--------------|--------------|
| 0e | fc | target | - | initiator | offline |
| 0f | fc | target | - | initiator | offline |
| 0g | fc | target | - | initiator | offline |
| 0h | fc | target | - | initiator | offline |
| 1a | fc | target | - | - | online |
| 1b | fc | target | - | - | online |

6 entries were displayed.

UTA/UTA2 ports can be configured into native FC mode or UTA/UTA2 mode. FC mode supports FC initiator and FC target; UTA/UTA2 mode allows concurrent NIC and FCoE traffic sharing the same 10GbE SFP+ interface and supports FC targets.

UTA/UTA2 ports might be found on an adapter or on the controller, and have the following configurations, but you should check the configuration of the UTA/UTA2 ports on the node3 and change it, if necessary:

- UTA/UTA2 cards ordered when the controller is ordered are configured before shipment to have the personality you request.
- UTA/UTA2 cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard UTA/UTA2 ports on new controllers are configured before shipment to have the personality you request.



Attention: If your system has storage disks, you enter the commands in this section at the cluster prompt unless directed to enter Maintenance mode. If you have a V- Series system or have FlexArray Virtualization Software and are connected to storage arrays, you enter commands in this section at the Maintenance mode prompt. You must be in Maintenance mode to configure UTA/UTA2 ports.

Steps

1. Check how the ports are currently configured by entering the following command on node3:

| If the system... | Then... |
|--|---------------------------|
| Has storage disks | No action required. |
| Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays | <code>ucadmin show</code> |

The system displays output similar to the following example:

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

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

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

3. Examine the output of the `ucadmin show` command and determine whether the UTA/UTA2 ports have the personality you want.
4. Take one of the following actions:

| If the UTA/UTA2 ports... | Then... |
|---|---|
| Do not have the personality that you want | Go to Step 5 . |
| Have the personality that you want | Skip Step 5 through Step 12 and go to Step 13 . |

5. Take one of the following actions:

| If you are configuring... | Then... |
|---------------------------|--|
| Ports on a UTA/UTA2 card | Go to Step 7 |
| Onboard UTA/UTA2 ports | Skip Step 7 and go to Step 8 . |

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 `cna`.
- `-t` is the FC4 type, `target` or `initiator`.



You must use FC initiator for tape drives, FlexArray Virtualization systems, and MetroCluster configurations. You must use the FC target for SAN clients.

8. Verify the settings:

```
ucadmin show
```

9. Verify the settings:

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

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
      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
*>
```

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

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

11. Cable the port.
12. Take one of the following actions:

| If the system... | Then... |
|--|--|
| Has storage disks | Go to Reassign node1 disks to node3 . |
| Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays | Return to <i>Install and boot node3</i> , and resume the procedure Step 23 . |

13. Exit maintenance mode:

```
halt
```

14. Boot the node into boot menu by running `boot_ontap` menu. If you are upgrading to an AFF A800 system, go to [Reassign node1 disks to node3, Step 9](#).

Reassign node1 disks to node3

You need to reassign the disks that belonged to node1 to node3 before verifying the node3 installation.

About this task

You perform the steps in this section on node3.

Steps

1. Go to the boot menu and using 22/7, select the hidden option `boot_after_controller_replacement`. At the prompt, enter node1 to reassign the disks of node1 to node3, as per the following example.

Expand the console output example

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

```

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

```

```

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

```



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

2. If the system goes into a reboot loop with the message `no disks found`, it indicates that the system has reset the FC or UTA/UTA2 ports back to the target mode and therefore is unable to see any disks. To resolve this continue with [Step 3](#) to [Step 8](#), or go to section [Verify the node3 installation](#).
3. Press Ctrl-C during AUTOBOOT to stop the node at the `LOADER>` prompt.
4. At the `LOADER` prompt, enter maintenance mode:

```
boot_ontap maint
```

5. In maintenance mode, display all the previously set initiator ports that are now in target mode:

```
ucadmin show
```

Change the ports back to initiator mode:

```
ucadmin modify -m fc -t initiator -f adapter name
```

6. Verify that the ports have been changed to initiator mode:

```
ucadmin show
```

7. Exit maintenance mode:

```
halt
```



If you are upgrading from a system that supports external disks to a system that also supports external disks, go to [Step 8](#).

If you are upgrading from a system that supports external disks to a system that supports both internal and external disks, for example, an AFF A800 system, go to [Step 9](#).

8. At the LOADER prompt, boot up:

```
boot_ontap menu
```

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

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



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

a. Return to the special boot menu:

```
LOADER> boot_ontap menu
```

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

Selection (1-11)? 10
```

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

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

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

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

e. Enter the backup data when prompted.



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

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



You might encounter an error at this stage. If an error occurs, repeat the substeps in [Step 8](#) until the system boots normally.

9. If you are upgrading from a system with external disks to a system that supports internal and external disks (AFF A800 systems, for example), set the node1 aggregate as the root aggregate to confirm that node3 boots from the root aggregate of node1. To set the root aggregate, go to the boot menu and select option 5 to enter maintenance mode.



You must perform the following substeps in the exact order shown; failure to do so might cause an outage or even data loss.

The following procedure sets node3 to boot from the root aggregate of node1:

- a. Enter maintenance mode:

```
boot_ontap maint
```

- b. Check the RAID, plex, and checksum information for the node1 aggregate:

```
aggr status -r
```

- c. Check the status of the node1 aggregate:

```
aggr status
```

- d. If necessary, bring the node1 aggregate online:

```
aggr_online root_aggr_from_node1
```

- e. Prevent the node3 from booting from its original root aggregate:

```
aggr offline root_aggr_on_node3
```

- f. Set the node1 root aggregate as the new root aggregate for node3:

```
aggr options aggr_from_node1 root
```

- g. Verify that the root aggregate of node3 is offline and the root aggregate for the disks brought over from node1 is online and set to root:

```
aggr status
```



Failing to perform the previous substep might cause node3 to boot from the internal root aggregate, or it might cause the system to assume a new cluster configuration exists or prompt you to identify one.

The following shows an example of the command output:

| Aggr | State | Status | Options |
|----------------------|---------|--|-----------------|
| aggr0_nst_fas8080_15 | online | raid_dp, aggr fast zeroed 64-bit | root, nosnap=on |
| aggr0 | offline | raid_dp, aggr fast zeroed 64-bit | diskroot |

Verify the node3 installation

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

About this task

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

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

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

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

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

Steps

1. Verify that node3 has joined quorum:

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

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

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

```
cluster show
```

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

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

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

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

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

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

Repeat Step 3 to verify that the cluster LIF is now listening on port 7700.

- Switch to advanced privilege mode:

```
set advanced
```

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

```
system controller replace show
```

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

- If you are working on a MetroCluster system, verify that the replaced controller is configured correctly for the MetroCluster configuration; the MetroCluster configuration should be in a healthy state. See [Verify the health of the MetroCluster configuration](#).

Reconfigure the intercluster LIFs on MetroCluster node node3, and check cluster peering to restore communication between the MetroCluster nodes before proceeding to Step 6.

Check the MetroCluster node status:

```
metrocluster node show
```

8. Resume the controller replacement operation:

```
system controller replace resume
```

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

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

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

2 entries were displayed.
```



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

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

Restore network configuration on node3

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

About this task

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



If you are changing the port speed of the e0a and e1a cluster ports on AFF A800 or AFF C800 systems, you might observe malformed packets being received after the speed conversion. See [NetApp Bugs Online Bug ID 1570339](#) and the knowledge base article [CRC errors on T6 ports after converting from 40GbE to 100GbE](#) for guidance.

Steps

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

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

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

2. List the broadcast domains on the cluster:

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

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

```
network port reachability show
```

You should see output like the following example:

```
clusterA::*> reachability show -node node1_node3
(network port reachability show)
Node          Port          Expected Reachability  Reachability Status
-----
node1_node3
a0a           Default:Default        no-reachability
a0a-822       Default:822           no-reachability
a0a-823       Default:823           no-reachability
e0M           Default:Mgmt           ok
e0a           Cluster:Cluster        misconfigured-
reachability
e0b           Cluster:Cluster        no-reachability
e0c           Cluster:Cluster        no-reachability
e0d           Cluster:Cluster        no-reachability
e0e           Cluster:Cluster        ok
e0e-822       -                      no-reachability
e0e-823       -                      no-reachability
e0f           Default:Default        no-reachability
e0f-822       Default:822           no-reachability
e0f-823       Default:823           no-reachability
e0g           Default:Default        misconfigured-
reachability
e0h           Default:Default        ok
e0h-822       Default:822           ok
e0h-823       Default:823           ok
18 entries were displayed.
```

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

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

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

You should see output like the following example:

```
Cluster ::> reachability repair -node node1_node3 -port e0h
```

```
Warning: Repairing port "node1_node3: e0h" may cause it to move into a
different broadcast domain, which can cause LIFs to be re-homed away
from the port. Are you sure you want to continue? {y|n}:
```

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

from the reachability status of the broadcast domain where it is currently located. Review the connectivity of the port and answer *y* or *n* as appropriate.

Verify that all physical ports have their expected reachability:

```
network port reachability show
```

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

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

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

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

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

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

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

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

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

6. Assign appropriate physical ports to the *Cluster* broadcast domain by using the following steps:

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

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

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

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

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

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

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

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

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

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

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

a. List displaced VLANs:

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

Output like the following should display:

```
Cluster::*> displaced-vlans show
(cluster controller-replacement network displaced-vlans show)
      Original
Node   Base Port   VLANs
-----
Node1   a0a           822, 823
        e0e           822, 823
2 entries were displayed.
```

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

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

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

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

The following is an example of restoring displaced VLANs on port "e0e" to 'e0h':

```
Cluster::*> displaced-vlans restore -node node1_node3 -port e0e
-destination-port e0h
```

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

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

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

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

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

10. Verify port reachability:

```
network port reachability show
```

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

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

```
network port show
```

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

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

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

- a. List any LIFs that are displaced:

```
displaced-interface show
```

- b. Restore LIF home nodes and home ports:

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

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

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

Restore key-manager configuration on node3

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

About this task

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

Steps

1. Run the following command from node3:

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

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

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

Example

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

| node | vserver | key-server | key-id |
|----------|---------|------------|---|
| restored | | | |
| ----- | ----- | ----- | ----- |
| node3 | svm1 | " " | 00000000000000000200000000000a008a81976 |
| true | | | 2190178f9350e071fbb90f00000000000000000 |

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

After you verify network configuration on node3 and before you relocate aggregates from node2 to node3, you must verify that the NAS data LIFs belonging to node1 that are currently on node2 are relocated from node2 to node3. You must also verify that the SAN LIFs exist on node3.

About this task

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



If you are changing the port speed of the T6-based Ethernet network interface cards or motherboard ports, you might observe malformed packets being received after the speed conversion. See [NetApp Bugs Online Bug ID 1570339](#) and the knowledge base article [CRC errors on T6 ports after converting from 40GbE to 100GbE](#) for guidance.

Steps

1. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

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

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

2. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following checks:

- Cluster health check
- Cluster LIF status check

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

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

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

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

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

4. If necessary, restore and revert any displaced LIFs. List any displaced LIFs:

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

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

```
cluster controller-replacement network displaced-interface restore-home-node
```

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

```
system controller replace resume
```

The system performs the following post-checks:

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

Stage 4. Relocate and retire node2

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

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

Before you begin

After the post-checks from the previous stage complete, the resource release for node2 starts automatically. The non-root aggregates and non-SAN data LIFs are migrated from node2 to node3.

About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. Moving SAN LIFs is not necessary for cluster or service health during the upgrade.

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



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

Steps

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

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

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

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

| Aggregate | Size | Available | Used% | State | #Vols | Nodes |
|-----------|---------|-----------|-------|--------|-------|-------|
| RAID | Status | | | | | |
| ----- | ----- | ----- | ----- | ----- | ----- | ----- |
| ----- | ----- | | | | | |
| aggr_1 | 744.9GB | 744.8GB | 0% | online | 5 | node2 |
| raid_dp | normal | | | | | |
| aggr_2 | 825.0GB | 825.0GB | 0% | online | 1 | node2 |
| raid_dp | normal | | | | | |

2 entries were displayed.

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

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

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

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

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

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

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

3. Verify that the LIFs have been moved to the correct ports and have a status of `up`. If any LIFs are down, set the administrative status of the LIFs to `up` by entering the following command, once for each LIF:

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

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

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

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

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

Retire node2

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

Steps

1. Resume the operation:

```
system controller replace resume
```

The node halts automatically.

After you finish

You can decommission node2 after the upgrade is completed. See [Decommission the old system](#).

Stage 5. Install and boot node4

Install and boot node4

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

About this task

The relocation operation is paused at the beginning of this stage. This process is mostly automated; the

operation pauses to enable you to check its status. You must manually resume the operation.

You need to netboot node4 if it does not have the same version of ONTAP 9 that is installed 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 V-Series system connected to storage arrays or a system with FlexArray Virtualization software that is connected to storage arrays, you must complete [Step 1](#) through [Step 21](#), then leave this section and follow instructions to [Configure FC ports on node4](#) and to [Check and configure UTA/UTA2 ports on node4](#), entering commands in Maintenance mode. You must then return to this section and resume with [Step 23](#).
- 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 instructions in the *Installation and Setup Instructions* or the *FlexArray Virtualization Installation Requirements and Reference* for the node4 platform, the appropriate disk shelf document, and *HA pair management*.

Refer to [References](#) to link to the *FlexArray Virtualization Installation Requirements and Reference* and *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.



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




Refer to [Prepare for netboot](#).

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



You must use the management port and IP as the netboot connection. Don't use a data LIF IP or a data outage might occur while the upgrade is being performed.

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

| If Dynamic Host Configuration Protocol (DHCP) is... | Then... |
|---|--|
| Not running | <p>Manually configure the connection by entering the following command at the boot environment prompt:</p> <pre>ifconfig e0M -addr=<i>filer_addr</i> -mask=<i>netmask</i> -gw=<i>gateway</i> -dns=<i>dns_addr</i> -domain=<i>dns_domain</i></pre> <p><i>filer_addr</i> is the IP address of the storage system (mandatory). <i>netmask</i> is the network mask of the storage system (mandatory). <i>gateway</i> is the gateway for the storage system (mandatory). <i>dns_addr</i> is the IP address of a name server on your network (optional). <i>dns_domain</i> is the DNS domain name (optional).</p> <div>  <p>Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p> </div> |

7. Perform netboot on node4:

| For... | Then... |
|----------------------------|--|
| FAS/AFF8000 series systems | <code>netboot http://<web_server_ip/path_to_web-accessible_directory>/netboot/kernel</code> |
| All other systems | <code>netboot http://<web_server_ip/path_to_web-accessible_directory>/<ontap_version>_image.tgz</code> |

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



Do not interrupt the boot.

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

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

Disregard the following message:

This procedure is not supported for Non-Disruptive Upgrade on an HA pair

The note applies to nondisruptive upgrades of ONTAP, and not upgrades of controllers.



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

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

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

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

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

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

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

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

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

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

```
ha-config show
```

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

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



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

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

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

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

If you have a MetroCluster configuration, use the following commands to modify the controller and chassis:

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

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

14. Exit maintenance mode:

```
halt
```

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

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

```
date
```

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

```
show date
```

17. If necessary, set the date on node4:

```
set date mm/dd/yyyy
```

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

```
show time
```

19. If necessary, set the time on node4:

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

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

```
setenv partner-sysid node3_sysid
```

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


Save the settings:

```
saveenv
```

21. Verify the partner-sysid for node4:

```
printenv partner-sysid
```

22. Take one of the following actions:

| If your system... | Then... |
|---|--|
| Has disks and no back-end storage | Go to Step 27 . |
| Is a V-Series system or a system with FlexArray Virtualization software connected to storage arrays | <div><div><div>a. Go to section Set the FC or UTA/UTA2 configuration on node4 and complete the subsections in this section.</div><div>b. Return to this section and complete the remaining steps, beginning with Step 23.</div></div><div><div></div><div>You must reconfigure FC onboard ports, CNA onboard ports, and CNA cards before you boot ONTAP on the V-Series or system with FlexArray Virtualization software.</div></div></div> |

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

If your system has a tape SAN, then you need zoning for the initiators. If required, modify the onboard ports to initiator by referring to the [Configure FC ports on node4](#). See your storage array and zoning documentation for further instructions on zoning.

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

See your storage array and zoning documentation for instructions.

25. Modify the worldwide port name (WWPN) values in the host or volume groups associated with array LUNs on the storage array.

Installing a new controller module changes the WWPN values associated with each onboard FC port.

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

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



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

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

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



You cannot mix FIPS drives with other types of drives on the same node or HA pair. You can mix SEDs with non-encrypting drives on the same node or HA pair.

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

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

28. Boot the node into boot menu:

```
boot_ontap menu
```

If you don't have an FC or UTA/UTA2 configuration, preform [Reassign node2 disks to node4, Step 1](#) so that node4 can recognize node2's disks.

29. For MetroCluster configuration, V-Series systems and systems with FlexArray Virtualization software connected to storage arrays you must set and configure the FC or UTA/UTA2 ports on node4 to detect the disks attached to the node.

To complete this task, go to section [Set the FC or UTA/UT2 configuration on node4](#).

Set the FC or UTA/UTA2 configuration on node4

If node4 has onboard FC ports, onboard unified target adapter (UTA/UTA2) ports, or a UTA/UTA2 card, you must configure the settings before completing the rest of the procedure.

About this task

You might need to complete the [Configure FC ports on node4](#) or [Check and configure UTA/UTA2 ports on node4](#) section, or both sections.



If node4 does not have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card, and you are upgrading a system with storage disks, you can skip to [Reassign node2 disks to node4](#).

However, if you have a V-Series system or have FlexArray Virtualization Software and are connected to storage arrays, and node4 does not have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card, you must return to *Install and boot node4* and resume the procedure at [Step 22](#).

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 FC adapter, you must set port configurations on the node before you bring it into service because the ports are not preconfigured. If the ports are not configured, you might experience a disruption in service.

Before you begin

You must have the values of the FC port settings from node2 that you saved in the section [Prepare the nodes for upgrade](#).

About this task

You can skip this section if your system does not have FC configurations. If your system has onboard UTA/UTA2 ports or a UTA/UTA2 adapter, you configure them in [Check and configure UTA/UTA2 ports on node4](#).



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

Steps


1. Take one of the following actions:

| If the system that you are upgrading... | Then... |
|---|--|
| Has storage disks | <code>system node hardware unified-connect show</code> |

| If the system that you are upgrading... | Then... |
|--|---------------------------|
| Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays | <code>ucadmin show</code> |

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

2. Compare the FC settings on node4 with the settings that you captured earlier from node1.
3. Take one of the following actions:

| If the system that you are upgrading... | Then... |
|--|--|
| Has storage disks | <p>Modify the FC ports on node4 as needed:</p> <ul style="list-style-type: none"> • To program target ports: <pre>ucadmin modify -m fc -t target <i>adapter</i></pre> <ul style="list-style-type: none"> • To program initiator ports: <pre>ucadmin modify -m fc -t initiator <i>adapter</i></pre> <p>-t is the FC4 type: target or initiator.</p> |
| Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays | <p>Modify the FC ports on node4 as needed:</p> <pre>ucadmin modify -m fc -t initiator -f <i>adapter_port_name</i></pre> <p>-t is the FC4 type, target or initiator.</p> <div>  <p>The FC ports must be programmed as initiators.</p> </div> |

4. Exit Maintenance mode:

```
halt
```

5. Boot the system from LOADER prompt:

```
boot_ontap menu
```

6. After you enter the command, wait until the system stops at the boot environment prompt.
7. Select option 5 from the boot menu for maintenance mode.
8. Take one of the following actions:

| If the system that you are upgrading... | Then... |
|--|--|
| Has storage disks | <ul style="list-style-type: none"> • Skip this section and go to Reassign node2 disks to node4 if node4 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports. |
| Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays | <ul style="list-style-type: none"> • Go to Check and configure UTA/UTA2 ports on node4 if node4 has a UTA/UTA2 card or UTA/UTA2 onboard ports. • Skip <i>Check and configure UTA/UTA2 ports on node4</i> if node4 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports, return to <i>Install and boot node4</i>, and resume the procedure at Step 23. |

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.



Attention: If your system has storage disks, you enter the commands in this section at the cluster prompt unless directed to enter Maintenance mode. If you have a MetroCluster FC system, V-Series system or a system with FlexArray Virtualization software that is connected to storage arrays, you must be in Maintenance mode to configure UTA/UTA2 ports.

Steps

1. Check how the ports are currently configured by using one of the following commands on node4:

| If the system... | Then... |
|-------------------|--|
| Has storage disks | <code>system node hardware unified-connect show</code> |

| If the system... | Then... |
|--|---------------------------|
| Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays | <code>ucadmin show</code> |

The system displays output similar to the following example:

```
*> ucadmin show
```

| Node | Adapter | Current Mode | Current Type | Pending Mode | Pending Type | Admin Status |
|------|---------|--------------|--------------|--------------|--------------|--------------|
| f-a | 0e | fc | initiator | - | - | online |
| f-a | 0f | fc | initiator | - | - | online |
| f-a | 0g | cna | target | - | - | online |
| f-a | 0h | cna | target | - | - | online |
| f-a | 0e | fc | initiator | - | - | online |
| f-a | 0f | fc | initiator | - | - | online |
| f-a | 0g | cna | target | - | - | online |
| f-a | 0h | cna | target | - | - | online |

```
*>
```

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

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

- Examine the output of the `ucadmin show` command and determine whether the UTA/UTA2 ports have the personality you want.
- Take one of the following actions:

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

- Take one of the following actions:

| If you are configuring... | Then... |
|---------------------------|--|
| Ports on a UTA/UTA2 card | Go to Step 7 |
| Onboard UTA/UTA2 ports | Skip Step 7 and go to Step 8 . |

- If the adapter is in initiator mode, and if the UTA/UTA2 port is online, take the UTA/UTA2 port offline:

```
storage disable adapter adapter_name
```

Adapters in target mode are automatically offline in Maintenance mode.

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, FlexArray Virtualization systems, and MetroCluster configurations. You must use the FC target for SAN clients.

8. Verify the settings by using the following command and examining its output:

```
ucadmin show
```

9. Verify the settings:

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

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.
*>
```

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

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

11. Cable the port.

12. Take one of the following actions:

| If the system... | Then... |
|--|---|
| Has storage disks | Go to Reassign node2 disks to node4 . |
| Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays | Return to <i>Install and boot node4</i> , and resume at Step 23 . |

13. Exit Maintenance mode:

```
halt
```

14. Boot the node into boot menu:

```
boot_ontap menu.
```

If you are upgrading to an AFF A800 system, go to [Reassign node 2 disks to node 4, Step 9](#).

Reassign node2 disks to node4

You need to reassign the disks that belonged to node2 to node4 before verifying the node4 installation..

About this task

You perform the steps in this section on node4.

Steps

1. Go to the boot menu and using 22/7, select the hidden option `boot_after_controller_replacement`. At the prompt, enter node2 to reassign the disks of node2 to node4, as per the following example.

Expand the console output example

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



```

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

```

```

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

```



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

2. If the system goes into a reboot loop with the message `no disks found`, it indicates that the system has reset the FC or UTA/UTA2 ports back to the target mode and therefore is unable to see any disks. To

resolve this, continue with [Step 3](#) to [Step 8](#) or go to section [Verify the node4 installation](#).

3. Press Ctrl-C during AUTOBOOT to stop the node at the LOADER> prompt.

4. At the LOADER prompt, enter maintenance mode:

```
boot_ontap maint
```

5. In maintenance mode, display all the previously set initiator ports that are now in target mode:

```
ucadmin show
```

Change the ports back to initiator mode:

```
ucadmin modify -m fc -t initiator -f adapter name
```

6. Verify that the ports have been changed to initiator mode:

```
ucadmin show
```

7. Exit maintenance mode:

```
halt
```



If you are upgrading from a system that supports external disks to a system that also supports external disks, go to [Step 8](#).

If you are upgrading from a system that uses external disks to a system that supports both internal and external disks, for example, an AFF A800 system, go to [Step 9](#).

8. At the LOADER prompt, boot up:

```
boot_ontap menu
```

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

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



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

a. Return to the special boot menu:

```
LOADER> boot_ontap menu
```

Please choose one of the following:

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

Selection (1-11)? 10

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

c. Enter `y` at the following prompt:

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

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

e. Enter the backup data when prompted.



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

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



You might encounter an error at this stage. If an error occurs, repeat the substeps in [Step 8](#) until the system boots normally.

9. If you are upgrading from a system with external disks to a system that supports internal and external disks (AFF A800 systems, for example), set the node2 aggregate as the root aggregate to ensure node4 boots from the root aggregate of node2. To set the root aggregate, go to the boot menu and select option 5 to enter maintenance mode.



You must perform the following substeps in the exact order shown; failure to do so might cause an outage or even data loss.

The following procedure sets node4 to boot from the root aggregate of node2:

a. Enter maintenance mode:

```
boot_ontap maint
```

b. Check the RAID, plex, and checksum information for the node2 aggregate:

```
aggr status -r
```

- c. Check the status of the node2 aggregate:

```
aggr status
```

- d. If necessary, bring the node2 aggregate online:

```
aggr_online root_aggr_from_node2
```

- e. Prevent the node4 from booting from its original root aggregate:

```
aggr offline root_aggr_on_node4
```

- f. Set the node2 root aggregate as the new root aggregate for node4:

```
aggr options aggr_from_node2 root
```

- g. Verify that the root aggregate of node4 is offline and the root aggregate for the disks brought over from node2 is online and set to root:

```
aggr status
```



Failing to perform the previous substep might cause node4 to boot from the internal root aggregate, or it might cause the system to assume a new cluster configuration exists or prompt you to identify one.

The following shows an example of the command output:

```
-----  
Aggr State                               Status                               Options  
aggr 0_nst_fas8080_15 online             raid_dp, aggr                       root, nosnap=on  
                                           fast zeroed  
                                           64-bit  
aggr0 offline                           raid_dp, aggr                       diskroot  
                                           fast zeroed`  
                                           64-bit  
-----
```

Verify the node4 installation

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

About this task

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

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

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

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

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

Steps

- 1. Verify that node4 has joined quorum:

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

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

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

```
cluster show
```

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

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

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

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

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

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

Repeat Step 3 to verify that the cluster LIF is now listening on port 7700.

5. Switch to advanced privilege mode:

```
set advanced
```

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

```
system controller replace show
```

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

7. If you are working on a MetroCluster system, verify that the replaced controller is configured correctly for the MetroCluster configuration; the MetroCluster configuration should be in a healthy state. Refer to [Verify the health of the MetroCluster configuration](#).

Reconfigure the intercluster LIFs on MetroCluster node node4, and check cluster peering to restore communication between the MetroCluster nodes before proceeding to [Step 6](#).

Check the MetroCluster node status:

```
metrocluster node show
```

8. Resume the controller replacement operation:

```
system controller replace resume
```

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

```

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

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

```



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

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

Restore network configuration on node4

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

About this task

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



If you are changing the port speed of the e0a and e1a cluster ports on AFF A800 or AFF C800 systems, you might observe malformed packets being received after the speed conversion. See [NetApp Bugs Online Bug ID 1570339](#) and the knowledge base article [CRC errors on T6 ports after converting from 40GbE to 100GbE](#) for guidance.

Steps

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

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

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

2. List the broadcast domains on the cluster:

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

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

```
network port reachability show
```

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

```
clusterA::*> reachability show -node node2_node4
(network port reachability show)
```

| Node | Port | Expected Reachability | Reachability Status |
|--------------|---------|-----------------------|---------------------|
| ----- | | | |
| node2_node4 | | | |
| | a0a | Default:Default | no-reachability |
| | a0a-822 | Default:822 | no-reachability |
| | a0a-823 | Default:823 | no-reachability |
| | e0M | Default:Mgmt | ok |
| | e0a | Cluster:Cluster | misconfigured- |
| reachability | e0b | Cluster:Cluster | no-reachability |
| | e0c | Cluster:Cluster | no-reachability |
| | e0d | Cluster:Cluster | no-reachability |
| | e0e | Cluster:Cluster | ok |
| | e0e-822 | - | no-reachability |
| | e0e-823 | - | no-reachability |
| | e0f | Default:Default | no-reachability |
| | e0f-822 | Default:822 | no-reachability |
| | e0f-823 | Default:823 | no-reachability |
| | e0g | Default:Default | misconfigured- |
| reachability | e0h | Default:Default | ok |
| | e0h-822 | Default:822 | ok |
| | e0h-823 | Default:823 | ok |

18 entries were displayed.

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

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

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

The output looks like the following example:

```
Cluster ::> reachability repair -node node2_node4 -port e0h
```

```
Warning: Repairing port "node2_node4: e0h" may cause it to move into a
different broadcast domain, which can cause LIFs to be re-homed away
from the port. Are you sure you want to continue? {y|n}:
```

A warning message, as shown above, is expected for ports with a reachability status that might be different from the reachability status of the broadcast domain where it is currently located.

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

Verify that all physical ports have their expected reachability:

```
network port reachability show
```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

a. List displaced VLANs:

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

Output like the following should display:

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

      Original
Node      Base Port      VLANs
-----
Node1     a0a            822, 823
          e0e            822, 823
```

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

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

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

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

The following is an example of restoring displaced VLANs on port "e0e" to "e0h":

```
Cluster::*> displaced-vlans restore -node node2_node4 -port e0e
-destination-port e0h
```

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

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

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

physical ports.

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

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

10. Verify port reachability:

```
network port reachability show
```

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

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

```
network port show
```

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

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

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

- a. List any LIFs that are displaced:

```
displaced-interface show
```

- b. Restore LIF home ports:

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

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

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

Restore key-manager configuration on node4

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

About this task

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

Steps

1. Run the following command from node4:

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

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

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

Example

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


| node | vserver | key-server | key-id |
|----------|---------|------------|---|
| restored | | | |
| ----- | ----- | ----- | ----- |
| node4 | svm1 | " " | 00000000000000000200000000000a008a81976 |
| true | | | 2190178f9350e071fbb90f00000000000000000 |

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

After you verify network configuration on node4 and before you relocate aggregates from node3 to node4, you must verify that the NAS data LIFs belonging to node2 that are currently on node3 are relocated from node3 to node4. You must also verify that the SAN LIFs exist on node4.

About this task

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



If you are changing the port speed of the T6-based Ethernet network interface cards or motherboard ports, you might observe malformed packets being received after the speed conversion. See [NetApp Bugs Online Bug ID 1570339](#) and the knowledge base article [CRC errors on T6 ports after converting from 40GbE to 100GbE](#) for guidance.

Steps

1. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

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

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

2. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following checks:

- Cluster health check
- Cluster LIF status check

After performing these checks, the system relocates the non-root aggregates and NAS data LIFs owned by node2 to the new controller, node4.

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

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

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

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

4. If necessary, restore and revert any displaced LIFs. List any displaced LIFs:

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

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

```
cluster controller-replacement network displaced-interface restore-home-node
```

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

```
system controller replace resume
```

The system performs the following post-checks:

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

Stage 6. Complete the upgrade

Manage authentication using KMIP servers

With ONTAP 9.8 or later, you can use Key Management Interoperability Protocol (KMIP) servers to manage authentication keys.

Steps

1. Add a new controller:

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

2. Add the key manager:

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

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

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

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

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

Confirm that the new controllers are set up correctly

To confirm correct setup, you must enable the HA pair. You must also verify that node3 and node4 can access each other's storage and that neither owns data LIFs belonging to other nodes on the cluster. In addition, you must confirm that node3 owns node1's aggregates and that node4 owns node2's aggregates, and that the volumes for both nodes are online.

Steps

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

```
storage failover show
```

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


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

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

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

```
cluster show
```

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

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

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

```
network interface show
```

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

```
network interface revert
```

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

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

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

7. Determine whether any volumes are offline:

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

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

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

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

9. Install new licenses for the new nodes by using the following command for each node:

```
system license add -license-code <license_code,license_code,license_code...>
```

The license-code parameter accepts a list of 28 upper-case alphabetic character keys. You can add one license at a time, or you can add multiple licenses at once, separating each license key by a comma.

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

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

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

- Delete all expired licenses:

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

- Delete all unused licenses:

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

- Delete a specific license from a cluster by using the following commands on the nodes:

```
system license delete -serial-number <node1_serial_number> -package *
```

```
system license delete -serial-number <node2_serial_number> -package *
```

The following output is displayed:

```
Warning: The following licenses will be removed:  
<list of each installed package>  
Do you want to continue? {y|n}: y
```

Enter `y` to remove all of the packages.

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

```
system license show
```

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

12. If self-encrypting drives are being used in the configuration and you have set the `kmip.init.maxwait` variable to `off` (for example, in [Install and boot node4, Step 27](#)), you must unset the variable:

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

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

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

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

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

After you finish

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

Set up Storage Encryption on the new controller module

If the replaced controller or the HA partner of the new controller uses Storage Encryption, you must configure the new controller module for Storage Encryption, including installing SSL certificates and setting up key management servers.

About this task

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

Steps

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

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

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

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

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

- b. Repeat the previous step for each listed key management server. You can link up to four key management servers.
- c. Verify the that the key management servers were added successfully:

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

3. On the new controller module, run the key management setup wizard to set up and install the key management servers.

You must install the same key management servers that are installed on the existing controller module.

- a. Launch the key management server setup wizard on the new node:

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

- b. Complete the steps in the wizard to configure key management servers.

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

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

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

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

About this task

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

Onboard Key Manager

Configure NVE or NAE using the Onboard Key Manager.

Steps

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

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

External Key Management

Configure NVE or NAE using External Key Management.

Steps

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

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

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

- a. Add the key management server:

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

- b. Repeat the previous step for each listed key management server. You can link up to four key management servers.

- c. Verify that the key management servers were added successfully:

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

3. On the new controller module, run the key management setup wizard to set up and install the key management servers.

You must install the same key management servers that are installed on the existing controller module.

- a. Launch the key management server setup wizard on the new node:

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

- b. Complete the steps in the wizard to configure key management servers.

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

```
security key-manager external restore
```

This command needs the OKM passphrase

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

After you finish

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

Decommission the old system

After upgrading, you can decommission the old system through the NetApp Support Site. Decommissioning the system tells NetApp that the system is no longer in operation and removes it from support databases.

Steps

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

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

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

4. Select **Go!**

A table displays cluster information, including the serial numbers.

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

Resume SnapMirror operations

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

Steps

1. Verify the SnapMirror status on the destination:

```
snapmirror show
```

2. Resume the SnapMirror relationship:

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

Troubleshoot

Aggregate relocation failures

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

Check for aggregate relocation failure

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

Steps

1. Enter the following command and examine the output:

```
storage aggregate relocation show
```

The `storage aggregate relocation show` command shows you which aggregates were successfully relocated and which ones were not, along with the causes of failure.

2. Check the console for any EMS messages.
3. Take one of the following actions:
 - Take the appropriate corrective action, depending on the output of the `storage aggregate relocation show` command and the output of the EMS message.
 - Force relocation of the aggregate or aggregates by using the `override-vetoes` option or the `override-destination-checks` option of the `storage aggregate relocation start` command.

For detailed information about the `storage aggregate relocation start`, `override-vetoes`, and `override-destination-checks` options, refer to [References](#) to link to the *ONTAP 9.8 Commands: Manual Page Reference*.

Aggregates originally on node1 are owned by node4 after completion of the upgrade

At the end of the upgrade procedure, node3 should be the new home node of aggregates that originally had node1 as the home node. You can relocate them after the upgrade.

About this task

Aggregates might fail to relocate correctly, having node1 as their home node instead of node3 under the following circumstances:

- During Stage 3, when aggregates are relocated from node2 to node3.
Some of the aggregates being relocated have node1 as their home node. For example, such an aggregate could be called `aggr_node_1`. If relocation of `aggr_node_1` fails during Stage 3, and relocation cannot be forced, then the aggregate will be left behind on node2.
- After Stage 4, when node2 is replaced with node4.
When node2 is replaced, `aggr_node_1` will come online with node4 as its home node instead of node3.

You can fix the incorrect ownership problem after Stage 6 once storage failover has been enabled by completing the following steps:

Steps

1. Enter the following command to get a list of aggregates:

```
storage aggregate show -nodes node4 -is-home true
```

To identify aggregates that were not correctly relocated, refer to the list of aggregates with the home owner of node1 that you obtained in the section [Prepare the nodes for upgrade](#) and compare it with output of the above command.

2. Compare the output of Step 1 with the output you captured for node1 in the section [Prepare the nodes for upgrade](#) and note any aggregates that were not correctly relocated.
3. Relocate the aggregates left behind on node4:

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

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

4. Verify that node3 is now the home owner of the aggregates:

```
storage aggregate show -aggregate aggr1,aggr2,aggr3... -fields home-name
```

aggr1,aggr2,aggr3... is the list of aggregates that had node1 as the original home owner.

Aggregates that do not have node3 as home owner can be relocated to node3 using the same relocation command in [Step 3](#).

Reboots, panics, or power cycles

The system might crash – reboot, panic or go through a power cycle – during different stages of the upgrade.

The solution to these problems depends on when they occur.

Reboots, panics, or power cycles during the pre-check phase

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

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

About this task

Takeover and giveback can proceed normally.

Steps

1. Check the console for EMS messages that the system might have issued and take the recommended corrective action.
2. Continue with the node-pair upgrade procedure.

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

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

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

About this task

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

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

Steps

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

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

Node2 does not take over but it is still serving data from all non-root aggregates.

Steps

1. Bring up node1.
2. Continue with the node-pair upgrade procedure.

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

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

About this task

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

Steps

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

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

Node1 does not take over.

Steps

1. Bring up node2.

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

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

Node2 crashes during the first verification phase with HA pair disabled

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

Steps

1. Bring up node2.

A client outage occurs for all aggregates while node2 is booting up.

2. Continue with the node-pair upgrade procedure.

Node3 crashes during the first verification phase with HA pair disabled

Node2 does not take over but it is still serving data from all non-root aggregates.

Steps

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

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

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

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

About this task

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

Steps

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

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

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

About this task

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

Steps

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

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

Node2 or node3 crashes during the post-check phase

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

Steps

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

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

Node3 crashes during the second resource-release phase

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

About this task

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

Steps

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

Node2 crashes during the second resource-release phase

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

About this task

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

Steps

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

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

Node3 crashes during the second verification phase

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

About this task

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

Steps

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

Node4 crashes during the second verification phase

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

About this task

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

Steps

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

Issues that can arise in multiple stages of the procedure

Some issues can occur during different stages of the procedure.

Unexpected "storage failover show" command output

During the procedure, if the node that hosts all data aggregates panics or is rebooted accidentally, you might see unexpected output for the `storage failover show` command before and after the reboot, panic, or power cycle.

About this task

You might see unexpected output from the `storage failover show` command in Stage 2, Stage 3, Stage 4, or Stage 5.

The following example shows the expected output of the `storage failover show` command if there are no reboots or panics on the node that hosts all the data aggregates:

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

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

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

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

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

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

Steps

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

LIF migration failure

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

Steps

1. Verify that the port MTU size is the same as that of the source node.

For example, if the cluster port MTU size is 9000 on the source node, it should be 9000 on the destination node.

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

References

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

- [Reference content](#)
- [Reference sites](#)

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. |
| Fabric-attached MetroCluster Installation and Configuration | Describes how to install and configure the MetroCluster hardware and software components in a fabric configuration. |
| FlexArray Virtualization Installation Requirements and Reference | Contains cabling instructions and other information for FlexArray Virtualization systems. |
| 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. |

| Content | Description |
|--|--|
| Network Management | Describes how to configure and manage physical and virtual network ports (VLANs and interface groups), LIFs, routing, and host-resolution services in clusters; optimize network traffic by load balancing; and monitor the cluster by using SNMP. |
| ONTAP 9.0 Commands: Manual Page Reference | Describes syntax and usage of supported ONTAP 9.0 commands. |
| ONTAP 9.1 Commands: Manual Page Reference | Describes syntax and usage of supported ONTAP 9.1 commands. |
| ONTAP 9.2 Commands: Manual Page Reference | Describes syntax and usage of supported ONTAP 9.2 commands. |
| ONTAP 9.3 Commands: Manual Page Reference | Describes syntax and usage of supported ONTAP 9.3 commands. |
| ONTAP 9.4 Commands: Manual Page Reference | Describes syntax and usage of supported ONTAP 9.4 commands. |
| ONTAP 9.5 Commands: Manual Page Reference | Describes syntax and usage of supported ONTAP 9.5 commands. |
| ONTAP 9.6 Commands: Manual Page Reference | Describes syntax and usage of supported ONTAP 9.6 commands. |
| ONTAP 9.7 Commands: Manual Page Reference | Describes syntax and usage of supported ONTAP 9.7 commands. |
| ONTAP 9.8 Commands: Manual Page Reference | Describes syntax and usage of supported ONTAP 9.8 commands. |
| ONTAP 9.9.1 Commands: Manual Page Reference | Describes syntax and usage of supported ONTAP 9.9.1 commands. |
| ONTAP 9.10.1 Commands: Manual Page Reference | Describes syntax and usage of supported ONTAP 9.10.1 commands. |
| SAN management with the CLI | Describes how to configure and manage LUNs, igroups, and targets using the iSCSI and FC protocols, and namespaces and subsystems using the NVMe/FC protocol. |
| SAN configuration reference | Contains information about FC and iSCSI topologies and wiring schemes. |
| Upgrade by moving volumes or storage | Describes how to quickly upgrade controller hardware in a cluster by moving storage or volumes. Also describes how to convert a supported model to a disk shelf. |
| Upgrade ONTAP | Contains instructions for downloading and upgrading ONTAP. |
| 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. |

| Content | Description |
|--|--|
| 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 © 2025 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.