



Use "system controller replace" commands to upgrade AFF A700 to AFF A900 or FAS9000 to FAS9500

Upgrade controllers

NetApp
December 01, 2022

Table of Contents

- Use "system controller replace" commands to upgrade AFF A700 to AFF A900 or FAS9000 to FAS9500 1
 - Overview 1
 - Automate the controller upgrade process 2
 - Decide whether to use the aggregate relocation procedure 2
 - Required tools and documentation 3
 - Guidelines for upgrading controllers 3
 - Overview of the ARL upgrade 4
 - Stage 1. Prepare for upgrade 6
 - Stage 2. Relocate resources and retire the AFF A700 or FAS9000 node1 12
 - Stage 3. Boot node1 with the AFF A900 or FAS9500 controller and NVRAM modules 23
 - Stage 4. Relocate resources from node2 and retire node2 34
 - Stage 5. Install the AFF A900 or FAS9500 NVRAM and controller modules on node2 36
 - Stage 6. Boot node2 with the AFF A900 or FAS9500 controller and NVRAM modules 41
 - Stage 7. Complete the upgrade 51
 - Troubleshoot 57
 - References 63

Use "system controller replace" commands to upgrade AFF A700 to AFF A900 or FAS9000 to FAS9500

Overview

This procedure describes how to upgrade the controller hardware using aggregate relocation (ARL) for the following system configurations:

Method	ONTAP version	Supported systems	Description
Using system controller replace commands	9.10.1 or later	AFF A700 to AFF A900	Converts an AFF A700 system to an AFF A900 system, keeping the A700 chassis
Using system controller replace commands	9.10.1 P3 or later	FAS9000 to FAS9500	Converts a FAS9000 system to a FAS9500 system, keeping the FAS9000 chassis.

You can nondisruptively upgrade controller hardware on a high availability (HA) pair by swapping the controller and the NVRAM modules. First, you migrate the non-root aggregates between the old controller nodes. After installation, you then migrate the non-root aggregates from the old controller nodes to the replacement controller nodes. The data hosted on the nodes that you are upgrading is accessible during the upgrade procedure.

About this task

During this controller upgrade procedure, you swap only the two field replaceable units (FRUs) on each node on the AFF A700 or FAS9000 system with the new FRUs.



You do not need to move, disconnect, or reconnect the I/O cards, data cables, disk shelves, and disks.

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

During the procedure, you upgrade the original controller hardware with the replacement controller hardware, relocating the ownership of non-root aggregates. You migrate aggregates multiple times from node to node to ensure that at least one node is serving data from the aggregates throughout the upgrade procedure. You also migrate data LIFs between nodes in the cluster as you proceed.



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

Important information

- This procedure is complex and assumes that you have advanced ONTAP administration skills. You also should read and understand the [Guidelines for upgrading controllers](#) and the [Overview of the ARL upgrade](#) sections before beginning the upgrade.
- This procedure assumes that the replacement controller hardware is new and has not been used in

another system. The steps required to prepare used controllers with the `wipeconfig` command are not included in this procedure. You must contact technical support if the replacement controller hardware was previously used as part of another ONTAP cluster or as a standalone single node system.

- You can use this procedure to upgrade the controller hardware in clusters with more than two nodes; however, you need to perform the procedure separately for each HA pair in the cluster.
- This procedure applies only to AFF A700 and FAS9000 systems. For all other controller models that need upgrading to an AFF A900 or FAS9500 system, refer to [References](#) to link to the *Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later* and the *Using Aggregate Relocation to Manually Upgrade Controller Hardware Running ONTAP 9.8 or Later* content.
- The AFF A900 and FAS9500 systems only support high-line power (200V to 240V). If your AFF A700 or FAS9000 system is running on low-line power (100V to 120V), you must convert the AFF A700 or FAS9000 input power before using this procedure.
- If you are upgrading from an AFF A700 or FAS9000 system with downtime, you can upgrade controller hardware by moving storage or contact technical support. Refer to [References](#) to link to *Upgrade by moving volumes or storage*.

Automate the controller upgrade process

During this controller upgrade, you replace an AFF A700 controller with an AFF A900 controller or a FAS9000 controller with a FAS9500 controller. This content provides the steps for the automated procedure, which uses automatic disk assignment and network port reachability checks to simplify the controller upgrade experience.

Decide whether to use the aggregate relocation procedure

This content describes how to upgrade storage controllers in an AFF A700 HA pair with new AFF A900 controllers or controllers in a FAS9000 HA pair with new FAS9500 controllers, while keeping all the existing data and disks. This is a complex procedure that should be used only by experienced administrators.

Use this procedure under the following circumstances:

- You have verified with your NetApp sales representative that you have received the AFF A900 or FAS9500 controller and NVRAM modules and the parts required for the upgrade.
- You are running ONTAP 9.10.1 or later for an AFF A900 controller.
- You are running ONTAP 9.10.1P3 or later for a FAS9500 controller.
- You do not want to add the new controllers as a new HA pair to the cluster and migrate the data by using volume moves.
- You are experienced in administering ONTAP and are comfortable with the risks of working in diagnostic privilege mode.

Do **NOT** use this procedure under the following circumstances:

- You are using FlexArray Virtualization Software on the AFF A700 or FAS9000 systems.
- You are using a shared switch for cluster-interconnect and Ethernet attached storage.

For upgrading Fabric MetroCluster or MetroCluster IP configurations, see [References](#) to link to the

MetroCluster Upgrade and Expansion content.



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

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

Refer to [References](#) to link to the *ONTAP 9 Documentation Center* where you can access ONTAP 9 product documentation.

Required tools and documentation

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

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

Guidelines for upgrading controllers

Using ARL to upgrade a pair of controllers in the following system configurations:

- Upgrade a pair of AFF A700 controllers to an AFF A900 system running ONTAP 9.10.1 or later
- Upgrade a pair of FAS9000 controllers to a FAS9500 system running ONTAP 9.10.1P3 or later

Supported upgrades for ARL

Controller upgrade from an AFF A700 to an AFF A900 or from a FAS9000 to a FAS9500 using ARL and the NVRAM module is the supported combination. If you have received a new AFF A900 or FAS9500 system as a complete system, including a new chassis, refer to [References](#) to link to the *Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later* content.



This procedure strictly applies to an AFF A700 to AFF A900 upgrade running ONTAP 9.10.1 or later and a FAS9000 to FAS9500 upgrade running 9.10.1P3 or later. Do NOT use this procedure to perform an upgrade between any other system combinations.

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.

Troubleshoot

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

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

Overview of the ARL upgrade

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

Upgrade the node pair

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

ARL upgrade sequence overview

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

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

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

Stage	Steps
Stage 1: Prepare for upgrade	<ol style="list-style-type: none">1. Verify the upgrade hardware2. Run prechecks to verify that you can perform the operation3. Correct aggregate ownership if a precheck fails4. Enter the cluster-base license keys5. Get an IP address for storage encryption6. Manage authentication by using a Key Management Operability Protocol (KMIP) server7. Manage storage encryption by using an onboard key manager (OKM)8. Quiesce the SnapMirror relationships (optional) <p>Aggregate ownership at the end of Stage 1:</p> <ul style="list-style-type: none">• Node1 is the home owner and current owner of the node1 aggregates• Node2 is the home owner and current owner of the node2 aggregates

Stage	Steps
<p>Stage 2: Relocate resources from node1 and retire AFF A700 or FAS9000 node1</p>	<ol style="list-style-type: none"> 1. Relocate non-root aggregates and NAS data LIFs owned by node1 to node2 2. Relocate failed or vetoed aggregates 3. Retire node1 4. Remove the AFF A700 or FAS9000 controller and NVRAM modules 5. Install the AFF A900 or FAS9500 NVRAM and controller modules on node1 6. Netboot node1 <p>Aggregate ownership at the end of Stage 2:</p> <ul style="list-style-type: none"> • Node2 is the current owner of node1 aggregates • Node2 is the home owner and current owner of node2 aggregates
<p>Stage 3: Boot node1 with the AFF A900 or FAS9500 controller and NVRAM modules</p>	<ol style="list-style-type: none"> 1. Boot node1 with the AFF A900 or FAS9500 controller and NVRAM modules 2. Verify that node1 is successfully installed 3. Restore node1 network configuration 4. Restore key-manager configuration on the upgraded node1 5. Move non-root aggregates and NAS data LIFs owned by node1 from node2 to the upgraded node1 <p>Aggregate ownership at the end of Stage 3:</p> <ul style="list-style-type: none"> • Upgraded AFF A900 or FAS9500 node1 is the home owner and current owner of node1 aggregates • Node2 is the home owner and current owner of node2 aggregates
<p>Stage 4: Relocate resources from node2 and retire node2</p>	<ol style="list-style-type: none"> 1. Relocate non-root aggregates and NAS data LIFs from node2 to node1 2. Retire node2 <p>Aggregate ownership at the end of Stage 4:</p> <ul style="list-style-type: none"> • Upgraded node1 is the home owner and current owner of aggregates that originally belonged to node1 • Upgraded node1 is the current owner of node2 aggregates

Stage	Steps
Stage 5: Install the AFF A900 or FAS9500 NVRAM and controller modules on node2	<ol style="list-style-type: none"> 1. Install the AFF A900 or FAS9500 NVRAM and controller modules on node2 2. Netboot node2 <p>Aggregate ownership at the end of Stage 5:</p> <ul style="list-style-type: none"> • Upgraded node1 is the home owner and current owner of the aggregates that originally belonged to node1. • Upgraded node2 is the home owner and current owner of aggregates that originally belonged to node2.
Stage 6: Boot node2 with the AFF A900 or FAS9500 controller and NVRAM modules	<ol style="list-style-type: none"> 1. Boot node2 with the AFF A900 or FAS9500 controller and NVRAM modules 2. Verify that node2 is correctly installed 3. Restore node2 network configuration 4. Move non-root aggregates and NAS data LIFs back to node2
Stage 7: Complete the upgrade	<ol style="list-style-type: none"> 1. Verify that the new controllers are set up correctly 2. Set up Storage Encryption on the new controller module 3. Set up NetApp Volume Encryption on the new controller module. 4. Decommission the old system. 5. Resume NetApp SnapMirror operations

Stage 1. Prepare for upgrade

Verify the upgrade hardware

Before starting the upgrade, you must verify that you have the correct hardware for an AFF A700 to AFF A900 or a FAS9000 to a FAS9500 upgrade. For each HA pair that you are upgrading, you must have two system controller modules and two NVRAM modules. If there are parts missing, contact technical support or your NetApp sales contact for assistance.

Prepare the nodes for upgrade

During Stage 1, you must prepare the nodes for the upgrade and run a series of prechecks. You must also prepare to netboot the new controllers.

Steps

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

```
set -privilege advanced
```

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

You will see output similar to the following:

Warning:

1. Current ONTAP version is 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. Select `y`. You will see the following output:

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

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

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

Precheck	Description
Image Version Check	Checks that the nodes being replaced run the same version of ONTAP. If the ONTAP image versions are different, the task fails. The new nodes must have the same version of ONTAP 9.x installed on them that is installed on the original nodes. If the new nodes have a different version of ONTAP installed, you need to netboot the new controllers after you install them. For instructions on how to upgrade ONTAP, refer to References to link to <i>Upgrade ONTAP</i> .
HA Status Check	Checks if both the nodes being replaced are in a high availability (HA) pair configuration. If storage failover is not enabled for the controllers, the task fails.
Aggregate Status Check	If the nodes being replaced own aggregates for which they are not the home owner, the task fails. The nodes should not own any non-local aggregates.
Disk Status Check	If any nodes being replaced have missing or failed disks, the task fails. If any disks are missing, refer to References to link to <i>Disk and aggregate management with the CLI</i> , <i>Logical storage management with the CLI</i> , and <i>High Availability management</i> to configure storage for the HA pair.
Data LIF Status Check	Checks if any of the nodes being replaced have non-local data LIFs. The nodes should not contain any data LIFs for which they are not the home owner. If one of the nodes contains non-local data LIFs, the task fails.
Cluster LIF Status	Checks whether the cluster LIFs are up for both nodes. If the cluster LIFs are down, the task fails.
ASUP Status Check	If AutoSupport notifications are not configured, the task fails. You must enable AutoSupport before beginning the controller replacement procedure.
CPU Utilization Check	Checks if the CPU utilization is more than 50% for any of the nodes being replaced. If the CPU usage is more than 50% for a considerable period of time, the task fails.
Aggregate Reconstruction Check	Checks if reconstruction is occurring on any data aggregates. If aggregate reconstruction is in progress, the task fails.
Node Affinity Job Check	Checks if any node affinity jobs are running. If node affinity jobs are running, the check fails.

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

You must run and save the output of the following commands individually:

- `vserver services name-service dns show`
- `network interface show -curr-node nodename -role cluster,intercluster,node-mgmt,cluster-mgmt,data`
- `network port show -node node_name -type physical`
- `service-processor show -node * -instance`
- `network fcp adapter show -node node_name`
- `network port ifgrp show`
- `system node show -instance -node node_name`
- `run -node node_name sysconfig`
- `storage aggregate show -node node_name`
- `volume show -node node_name`
- `storage array config show -switch switch_name`
- `system license show -owner node_name`
- `storage encryption disk show`
- `security key-manager onboard show-backup`
- `security key-manager external show`
- `security key-manager external show-status`
- `reachability show -detail`



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

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

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

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

Refer to [References](#) to link to the *NetApp Support Site* where you can obtain new 2-character license keys for 9.10.1 or later. The keys are available in the *My Support* section under *Software licenses*. If the site does not have the license keys you need, contact your NetApp sales representative.

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

Get an IP address of an external key management server for Storage Encryption

After upgrading, you must immediately configure Storage Encryption and establish a cluster-wide authentication key to replace the previous node-level authentication keys.

Steps

1. Install the necessary client and server secure sockets layer (SSL) certificates required to communicate with key management servers:

```
security certificate install
```

2. Configure Storage Encryption on all nodes by using the following command on each node:

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

3. Add the IP address for each key management server:

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

4. Verify that the same key management servers are configured and available on all nodes in the cluster:

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

5. Create a new cluster-wide authentication key:

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

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

7. Rekey all self-encrypting drives with the new authentication key:

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

Manage authentication using KMIP servers

Beginning with ONTAP 9.10.1, you can use Key Management Interoperability Protocol (KMIP) servers to manage authentication keys.

Steps

1. Add a new controller:

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

2. Add the key manager:

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

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

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

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

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

5. Rekey all self-encrypting disks with the new authentication key:

```
storage encryption disk modify -disk * [-data-key-id nonMSID AK]
```

6. If you use the Federal Information Processing Standard (FIPS), rekey all self-encrypting disks with the new authentication key:

```
storage encryption disk modify -disk * [-fips-key-id nonMSID AK]
```

Manage storage encryption using Onboard Key Manager

You can use the OKM to manage encryption keys. If you plan to use OKM, you must record the passphrase and backup material before beginning the upgrade.

Steps

1. Save the passphrase to a secure location.
2. Create a backup for recovery purposes. Run the following command and save the output:

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

Quiesce the SnapMirror relationships (optional)

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

Steps

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

```
snapmirror show
```



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

```
snapmirror abort -destination-vserver vserver_name
```

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

2. Quiesce all relationships between the cluster:

```
snapmirror quiesce -destination-vserver vserver_name
```

Stage 2. Relocate resources and retire the AFF A700 or FAS9000 node1

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

Before you can replace node1 with the AFF A900 or FAS9500 controller and NVRAM modules, you must move the non-root aggregates and NAS data LIFs from node1 to node2 before eventually restoring node1's resources back on node1 running on the AFF A900 or FAS9500 system. This process is largely automated; the operation pauses to enable you to check its status.

Before you begin

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

About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. You are not required to move SAN LIFs for cluster or service health during the upgrade. You must verify that the LIFs are healthy and located on appropriate ports after you bring node1 online as the AFF A900 or FAS9500 system.



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

Steps

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

```
system controller replace resume
```

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

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

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

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

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

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

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

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

Aggregate  Size      Available  Used%  State  #Vols  Nodes  RAID Status
-----
-----
aggr_1     744.9GB  744.8GB   0%     online  5     node2
raid_dp,normal
aggr_2     825.0GB  825.0GB   0%     online  1     node2
raid_dp,normal
2 entries were displayed.
```

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

```
storage aggregate online -aggregate aggr_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.

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

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

Relocate failed or vetoed aggregates

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

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

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

- When prompted, enter `y`.
- 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, resume the automated operation to disable the HA pair with node2 and shut down node1 correctly. You must later remove the AFF A700 or the FAS9000 controller and NVRAM modules from the node1 chassis and then install the AFF A900 or the FAS9500 controller and NVRAM modules on node1.

Steps

- Resume the operation:

```
system controller replace resume
```

2. Verify that node1 has been halted:

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

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

Remove the AFF A700 or FAS9000 controller and NVRAM modules

At this stage, node1 is down and all data is served by node2. Because node1 and node2 are in the same chassis and powered by the same set of power supplies, do NOT power off the chassis. You must take care to remove only the node1 controller module and the node1 NVRAM module. Typically, node1 is controller A, located on the left side of the chassis when looking at the controllers from the rear of the system. The controller label is located on the chassis directly above the controller module.

Before you begin

If you are not already grounded, correctly ground yourself.

Remove the AFF A700 or FAS9000 controller module

Use the following procedure to remove the AFF A700 or FAS9000 controller module.


Steps

1. Detach the console cable, if any, and the management cable from the node1 controller module before removing the controller module from node1.



When you are working on node1, you only remove the console and e0M cables from node1. You must not remove or change any other cables or connections on either node1 or node2 during this process.

2. Unlock and remove the controller module A from the chassis.
 - a. Slide the orange button on the cam handle downward until it unlocks.



1	Cam handle release button
2	Cam handle

- b. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

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

Remove the AFF A700 or FAS9000 NVRAM module

Use the following procedure to remove the AFF A700 or FAS9000 NVRAM module.




The AFF A700 or FAS9000 NVRAM module is in slot 6 and is double the height of the other modules in the system.

Steps

1. Unlock and remove the NVRAM module from slot 6 of node1.
 - a. Depress the lettered and numbered cam button.

The cam button moves away from the chassis.
 - b. Rotate the cam latch down until it is in a horizontal position.

The NVRAM module disengages from the chassis and moves a few inches.
 - c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module face.



1	Lettered and numbered I/O cam latch
----------	-------------------------------------

2	I/O latch completely unlocked
----------	-------------------------------

Install the AFF A900 or FAS9500 NVRAM and controller modules on node1

Install the AFF A900 or FAS9500 NVRAM and controller modules that you received for the upgrade on node1.

You must note the following when performing the installation:

- Move all blank filler modules in slots 6-1 and 6-2 from the old NVRAM module to the new NVRAM module.
- Do NOT move the coredump device from the AFF A700 NVRAM module to the AFF A900 NVRAM module.
- Move all flash cache modules installed in the FAS9000 NVRAM module to the FAS9500 NVRAM module.

Before you begin


If you are not already grounded, correctly ground yourself.

Install the AFF A900 or FAS9500 NVRAM module

Use the following procedure to install the AFF A900 or FAS9500 NVRAM module in slot 6 of node1.

Steps

1. Align the NVRAM module with the edges of the chassis opening in slot 6.
2. Gently slide the NVRAM module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the NVRAM module in place.



1	Lettered and numbered I/O cam latch
2	I/O latch completely unlocked

Install the AFF A900 or FAS9500 controller module on node1.

Use the following procedure to install the AFF A900 or FAS9500 controller module in node1.

Steps

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



Do not completely insert the controller module in the chassis until instructed to do so later in the procedure.

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




Because the chassis is already powered ON, node1 starts BIOS initialization followed by autoboot as soon as it is fully seated. To interrupt the node1 boot, before completely inserting the controller module into the slot, it is recommended that you connect the serial console and management cables to the node1 controller module.

3. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latch rises when the controller module is fully seated.



To avoid damaging the connectors, do not use excessive force when sliding the controller module into the chassis.



1	Cam handle locking latch
2	Cam handle in the unlocked position

4. Connect the serial console as soon as the module is seated and be ready to interrupt AUTOBOOT of

node1.

5. After you interrupt AUTOBOOT, node1 stops at the LOADER prompt. If you do not interrupt AUTOBOOT on time and node1 starts booting, wait for the prompt and press **Ctrl-C** to go into the boot menu. After the node stops at the boot menu, use option 8 to reboot the node and interrupt the AUTOBOOT during reboot.
6. At the LOADER> prompt of node1, set the default environment variables:

```
set-defaults
```

7. Save the default environment variables settings:

```
saveenv
```

Netboot node1

After swapping the corresponding AFF A900 or FAS9500 node1 controller and NVRAM modules, you must netboot node1. The term netboot means that you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you must add a copy of the ONTAP 9 boot image onto a web server that the system can access.

It is not possible to check the version of ONTAP installed on the boot media of an AFF A900 or FAS9500 controller module unless it is installed in a chassis and powered ON. The ONTAP version on the AFF A900 or FAS9500 boot media must be same as the ONTAP version running on the AFF A700 or FAS9000 system that you are upgrading and both the primary and backup boot images must match. You can configure the images by performing a netboot followed by the `wipeconfig` command from the boot menu. If the controller module was previously used in another cluster, the `wipeconfig` command clears any residual configuration on the boot media.



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

Before you begin

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

About this task

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

Steps

1. Refer to [References](#) to link to the *NetApp Support Site* to download the files used for performing the netboot of the system.
2. Download the appropriate ONTAP software from the software download section of the *NetApp Support Site* and store the `<ontap_version>_image.tgz` file on a web-accessible directory.
3. Change to the web-accessible directory and verify that the files you need are available.
4. Your directory listing should contain `<ontap_version>_image.tgz`.

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



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

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

6. Perform netboot on node1:

```
netboot http://<web_server_ip/path_to_web_accessible_directory>/netboot/kernel
```



Do not interrupt the boot.

7. Wait for the node1 now running on the AFF A900 or FAS9500 controller module to boot and display the boot menu options as shown below:

Please choose one of the following:

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

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

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



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

Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the wrong image might install. This issue applies to all ONTAP releases.

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

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

The `<path_to_the_web-accessible_directory>` should lead to where you downloaded the `<ontap_version>_image.tgz` in [Step 2](#).

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

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

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

- b. Enter `y` to reboot when you see the following prompt:

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

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

11. At the prompt, run the `wipeconfig` command to clear any previous configuration on the boot media:
 - a. When you see the message below, answer `yes`:

```
This will delete critical system configuration, including cluster
membership.
Warning: do not run this option on a HA node that has been taken
over.
Are you sure you want to continue?:
```

- b. The node reboots to finish the `wipeconfig` and then stops at the boot menu.
12. Select option 5 to go to maintenance mode from the boot menu. Answer `yes` to the prompts until the node stops at maintenance mode and the command prompt `*>`.
13. Verify that the controller and chassis are configured as `ha`:

```
ha-config show
```

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

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

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

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

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

15. Verify the `ha-config` settings:

```
ha-config show
```

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

16. Halt node1:

```
halt
```

Node1 should stop at the `LOADER` prompt.

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

date

18. On node1, check the date by using the following command at the boot environment prompt:

```
show date
```

19. If necessary, set the date on node1:

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



Set the corresponding UTC date on node1.

20. On node1, check the time by using the following command at the boot environment prompt:

```
show time
```

21. If necessary, set the time on node1:

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



Set the corresponding UTC time on node1.

22. Set the partner system ID on node1:

```
setenv partner-sysid node2_sysid
```

You can obtain the node2 system ID from the `node show -node node2` command output on node2.

- a. Save the settings:

```
saveenv
```

23. On node1, at the LOADER prompt, verify the `partner-sysid` for node1:

```
printenv partner-sysid
```

For node1, the `partner-sysid` must be that of node2.

Stage 3. Boot node1 with the AFF A900 or FAS9500 controller and NVRAM modules

Boot node1 with the AFF A900 or FAS9500 controller and NVRAM modules

Node1 with the AFF A900 or FAS9500 controller and NVRAM modules is now ready for boot up. An AFF A700 to AFF A900 or a FAS9000 to FAS9500 upgrade by swapping the controller and NVRAM modules involves moving only the console and management connections. This section provides the steps required to boot node1 with the AFF A900 or FAS9500 controller and NVRAM modules.

Steps

1. If NetApp Storage Encryption (NSE) is in use on this configuration, the `setenv bootarg.storageencryption.support` command must be set to `true`, and the `kmip.init.maxwait` variable must be set to `off` to avoid a boot loop after the node1 configuration is loaded:

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

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

2. Boot the node into the boot menu:

```
boot_ontap menu
```

3. Reassign the AFF A700 node1 disks to AFF A900 node1 by entering "22/7" and selecting the hidden option `boot_after_controller_replacement` when the node stops at the boot menu.

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

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

```
If [localhost:disk.encryptNoSupport:ALERT]: Detected FIPS-certified
encrypting drive and, or, [localhost:diskown.errorDuringIO:error]:
error 3 (disk failed) on disk errors occur, perform the following steps:
```



1. Halt the node at the LOADER prompt.
2. Check and reset the storage encryption bootargs mentioned in [Step 1](#).
3. At the loader prompt, boot up:

```
boot_ontap
```

You can use the following example as a reference:

Expand the console output example

```
LOADER-A> boot_ontap menu
.
.
<output truncated>
.
All rights reserved.
*****
*                                     *
* Press Ctrl-C for Boot Menu. *
*                                     *
*****
.
<output truncated>
.
Please choose one of the following:

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 22/7

(22/7)                                     Print this secret List
(25/6)                                     Force boot with multiple filesystem
disks missing.
(25/7)                                     Boot w/ disk labels forced to clean.
(29/7)                                     Bypass media errors.
(44/4a)                                    Zero disks if needed and create new
flexible root volume.
(44/7)                                     Assign all disks, Initialize all
disks as SPARE, write DDR labels
.
.
<output truncated>
.
.
(wipeconfig)                               Clean all configuration on boot
```

```
device
(boot_after_controller_replacement) Boot after controller upgrade
(boot_after_mcc_transition)          Boot after MCC transition
(9a)                                  Unpartition all disks and remove
their ownership information.
(9b)                                  Clean configuration and
initialize node with partitioned disks.
(9c)                                  Clean configuration and
initialize node with whole disks.
(9d)                                  Reboot the node.
(9e)                                  Return to main boot menu.
```

The boot device has changed. System configuration information could be lost. Use option (6) to restore the system configuration, or option (4) to initialize all disks and setup a new system. Normal Boot is prohibited.

Please choose one of the following:

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

This will replace all flash-based configuration with the last backup to disks. Are you sure you want to continue?: yes

.
.

<output truncated>

.
.

Controller Replacement: Provide name of the node you would like to replace:<nodename of the node being replaced>

Changing sysid of node nodel disks.

Fetches sanown old_owner_sysid = 536940063 and calculated old sys id

```
= 536940063
Partner sysid = 4294967295, owner sysid = 536940063
.
.
<output truncated>
.
.
varfs_backup_restore: restore using /mroot/etc/varfs.tgz
varfs_backup_restore: attempting to restore /var/kmip to the boot
device
varfs_backup_restore: failed to restore /var/kmip to the boot device
varfs_backup_restore: attempting to restore env file to the boot
device
varfs_backup_restore: successfully restored env file to the boot
device wrote key file "/tmp/rndc.key"
varfs_backup_restore: timeout waiting for login
varfs_backup_restore: Rebooting to load the new varfs
Terminated
<node reboots>

System rebooting...

.
.
Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...
Rebooting to load the restored env file...

.
System rebooting...

.
.
.
<output truncated>
.
.
.
.
WARNING: System ID mismatch. This usually occurs when replacing a
boot device or NVRAM cards!
Override system ID? {y|n} y
.
.
.
.
Login:
```



The system IDs shown in the preceding example are example IDs. The actual system IDs of the nodes that you are upgrading will be different.

Between entering node names at the prompt and the login prompt, the node reboots a few times to restore the environment variables, update firmware on the cards in the system, and for other ONTAP updates.

Verify the node1 installation

You must verify the node1 installation with the AFF A900 or FAS9500 controller and NVRAM modules. Because there is no change to physical ports, you are not required to map the physical ports from the old node1 to the replacement node1.

About this task

After you boot node1 with the AFF A900 or FAS9500 controller module, you must verify that it is installed correctly. You must wait for node1 to join quorum and then resume the controller replacement operation.

At this point in the procedure, the controller upgrade operation should have paused as node1 attempts to join quorum automatically.

Steps

1. Verify that node1 has joined quorum:

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

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

2. Verify that node1 is part of the same cluster as node2 and that it is healthy:

```
cluster show
```

3. Switch to advanced privilege mode:

```
set advanced
```

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

```
system controller replace show
```

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

5. Resume the controller replacement operation:

```
system controller replace resume
```

6. The controller replacement operation pauses for intervention with the following message:

```

Cluster::*> system controller replace show
Node           Status           Error-Action
-----
Node1          Paused-for-intervention  Follow the instructions given
in
Node2          None              Step Details

```

Step Details:

To complete the Network Reachability task, the ONTAP network configuration must be manually adjusted to match the new physical network configuration of the hardware. This includes:

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



In this procedure, section *Re-creating VLANs, ifgrps, and broadcast domains* has been renamed *Restore network configuration on node1*.

7. With the controller replacement in a paused state, proceed to [Restore network configuration on node1](#).

Restore network configuration on node1

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

About this task

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

Steps

1. List all the physical ports that are on upgraded node1:

```
network port show -node node1
```

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:

```
broadcast-domain show
```

3. List the network port reachability of all ports on node1:

```
network port reachability show -node node1
```

You should see output like the following example:

```
Cluster::> reachability show -node node1
(network port reachability show)
Node      Port      Expected Reachability      Reachability
Status
-----
Node1
    a0a      Default:Default      ok
    a0a-822  Default:822          ok
    a0a-823  Default:823          ok
    e0M      Default:Mgmt         ok
    e11a     -                    no-reachability
    e11b     -                    no-reachability
    e11c     -                    no-reachability
    e11d     -                    no-reachability
    e3a      -                    no-reachability
    e3b      -                    no-reachability
    e4a      Cluster:Cluster      ok
    e4e      Cluster:Cluster      ok
    e5a      -                    no-reachability
    e7a      -                    no-reachability
    e9a      Default:Default      ok
    e9a-822  Default:822          ok
    e9a-823  Default:823          ok
    e9b      Default:Default      ok
    e9b-822  Default:822          ok
    e9b-823  Default:823          ok
    e9c      Default:Default      ok
    e9d      Default:Default      ok
22 entries were displayed.
```

In the preceding example, node1 booted after the controller replacement. Some ports do not have reachability because there is no physical connectivity. You must repair any ports with a reachability status

other than `ok`.



During an AFF A700 to AFF A900 or a FAS9000 to FAS9500 controller upgrade, the network ports and their connectivity should not change. All ports should reside in the correct broadcast domains and the network port reachability should not change. However, before moving LIFs from node2 back to node1, you must verify the reachability and health status of the network ports.

4. Repair the reachability for each of the ports on node1 with a reachability status other than `ok` by using the following command, in the following order:

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

- a. Physical ports
- b. VLAN ports

You should see output like the following example:

```
Cluster ::> reachability repair -node node1 -port e11b
```

```
Warning: Repairing port "node1:e11b" may cause it to move into a  
different broadcast domain, which can cause LIFs to be re-homed away  
from the port. Are you sure you want to continue? {y|n}:
```

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

Verify that all physical ports have their expected reachability:

```
network port reachability show
```

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

5. Verify port reachability:

```
network port reachability show
```

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

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

```
network port show
```

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

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

8. Restore LIF home ports, specifying the Vserver and LIF home ports, if any, that need to be restored by using the following steps:
 - a. List any LIFs that are displaced:

```
displaced-interface show
```

- b. Restore LIF home nodes and home ports:

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

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

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

Restore key-manager configuration on the upgraded node1

If you are using NetApp Volume Encryption (NVE) to encrypt volumes on the system you are upgrading, the encryption configuration must be synchronized to the new nodes. If you do not restore key-manager, when you relocate the node1 aggregates from node2 to the upgraded node1 by using ARL, encrypted volumes will be taken offline.

About this task

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

Steps

1. Synchronize the encryption configuration for Onboard Key Manager (OKM) by using the following command at the cluster prompt:

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

2. Enter the cluster-wide passphrase for the OKM.

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

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

About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. Moving SAN LIFs is not necessary for cluster or service health during the upgrade. SAN LIFs are not moved unless they need to be mapped to new ports. After you bring node1 online, you must verify that the LIFs are healthy and located on the appropriate

ports.

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. Perform a network reachability check:

```
network port reachability -show-detail -node node1
```

Confirm that all connected ports, including the interface group and VLAN ports, show their status as OK.

3. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following checks:

- Cluster health check
- Cluster LIF status check

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

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

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

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

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

5. If necessary, restore and revert 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 node1:

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

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

```
system controller replace resume
```

The system performs the following post-checks:

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

Stage 4. Relocate resources from node2 and retire node2

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

Before you can replace node2 with the AFF A900 or the FAS9500 controller module and the NVRAM module, you must first relocate the non-root aggregates that are owned by node2 to node1.

Before you begin

After the post-checks from the previous stage complete, the resource release for node2 starts automatically. The non-root aggregates and non-SAN data LIFs are migrated from node2 to the new node1.

About this task

After the aggregates and LIFs are migrated, the operation is paused for verification purposes. At this stage, you must verify that all the non-root aggregates and non-SAN data LIFs are migrated to the new node1.

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

Steps

1. Verify that all the non-root aggregates are online and their state on node1:

```
storage aggregate show -node node1 -state online -root false
```

The following example shows that the non-root aggregates on node1 are online:

```
cluster::> storage aggregate show -node node1 state online -root false
```

Aggregate RAID	Size Status	Available	Used%	State	#Vols	Nodes
aggr_1	744.9GB	744.8GB	0%	online	5	node1
normal						raid_dp
aggr_2	825.0GB	825.0GB	0%	online	1	node1
normal						raid_dp

2 entries were displayed.

If the aggregates have gone offline or become foreign on node1, bring them online by using the following command on the new node1, once for each aggregate:

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

2. Verify that all the volumes are online on node1 by using the following command on node1 and examining its output:

```
volume show -node node1 -state offline
```

If any volumes are offline on node1, bring them online by using the following command on node1, once for each volume:

```
volume online -vserver 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  
nodename - status-admin up
```

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

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

Retire node2

To retire node2, you must shut node2 down correctly and 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 the AFF A900 or FAS9500 NVRAM and controller modules on node2

Install the AFF A900 or FAS9500 NVRAM and controller modules on node2

Install the AFF A900 or FAS9500 NVRAM and controller modules that you received for the upgrade on node2. Node2 is controller B located on the right side of the chassis when looking at the controllers from the rear of the system.

You must note the following when performing the installation:

- Move all blank filler modules in slots 6-1 and 6-2 from the old NVRAM module to the new NVRAM module.
- Do NOT move the coredump device from the AFF A700 NVRAM module to the AFF A900 NVRAM module.
- Move all flash cache modules installed in the FAS9000 NVRAM module to the FAS9500 NVRAM module.

Before you begin

If you are not already grounded, correctly ground yourself.

Install the AFF A900 or FAS9500 NVRAM module

Use the following procedure to install the AFF A900 or FAS9500 NVRAM module in slot 6 of node2.

Steps

1. Align the NVRAM module with the edges of the chassis opening in slot 6.
2. Gently slide the NVRAM module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the NVRAM module in place.

Install the AFF A900 or FAS9500 controller module in node2

Use the following procedure to install the AFF A900 or FAS9500 controller module in node2.

Steps

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



The bay label is located on the chassis directly above the controller module.



Do not completely insert the controller module in the chassis until you are instructed to do so later in the procedure.

2. Cable the management and console ports to the node2 controller module.



Because the chassis is already powered ON, node2 starts booting as soon as it is fully seated. To avoid node2 booting, it is recommended that you connect the console and management cables to the node2 controller module before completely inserting the controller module into the slot.

3. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latch rises when the controller module is fully seated.



To avoid damaging the connectors, do not use excessive force when sliding the controller module into the chassis.

4. Connect the serial console as soon as the module is seated and be ready to interrupt AUTOBOOT of node1.
5. After you interrupt AUTOBOOT, node2 stops at the LOADER prompt. If you do not interrupt AUTOBOOT on time and node2 starts booting, wait for the prompt and press **Ctrl-C** to go into the boot menu. After the node stops at the boot menu, use option 8 to reboot the node and interrupt the AUTOBOOT during reboot.

Netboot node2

After swapping the corresponding AFF A900 or FAS9500 node2 controller module and NVRAM, module you might need to netboot them. The term netboot means that you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you must put a copy of the ONTAP 9 boot image onto a web server that the system can access.

It is not possible to check the version of ONTAP installed on the boot media of an AFF A900 or a FAS9500 controller module unless it is installed in a chassis and powered ON. The ONTAP version on the AFF A900 or the FAS9500 boot media must be the same as the ONTAP version running on the AFF A700 or the FAS9000 system that you are upgrading and both the primary and backup boot images must match. You can configure the images by performing a netboot followed by the `wipeconfig` command from the boot menu. If the controller module was previously used in another cluster, the `wipeconfig` command clears any residual configuration on the boot media.



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

Before you begin

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

About this task

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

Steps

1. Refer to [References](#) to link to the *NetApp Support Site* to download the files used for performing the netboot of the system.
2. Download the appropriate ONTAP software from the software download section of the NetApp Support Site and store the `<ontap_version>_image.tgz` file on a web-accessible directory.
3. Change to the web-accessible directory and verify that the files you need are available.
4. Your directory listing should contain `<ontap_version>_image.tgz`.
5. Configure the netboot connection by choosing one of the following actions.



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

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

6. Perform netboot on node2:

```
netboot http://<web_server_ip/path_to_web_accessible_directory>/netboot/kernel
```



Do not interrupt the boot.

7. Wait for the node2 now running on the AFF A900 or FAS9500 controller module to boot and display the boot menu options as shown in the following output:

Please choose one of the following:

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

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

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



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

Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the wrong image might install. This issue applies to all ONTAP releases.

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

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

The `<path_to_the_web-accessible_directory>` should lead to where you downloaded the `<ontap_version>_image.tgz` in [Step 2](#).

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

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

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

- b. Enter `y` to reboot when you see the following prompt:

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

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

11. At the prompt, run the `wipeconfig` command to clear any previous configuration on the boot media.
 - a. When you see the message below, answer `yes`:

```
This will delete critical system configuration, including cluster
membership.
Warning: do not run this option on a HA node that has been taken
over.
Are you sure you want to continue?:
```

- b. The node reboots to finish the `wipeconfig` and then stops at the boot menu.
12. Select maintenance mode 5 from the boot menu and enter `y` when you are prompted to continue with the boot.
13. Verify that the controller and chassis are configured as `ha`:

```
ha-config show
```

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

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

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

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

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

15. Halt node2:

```
halt
```

Node2 should stop at the `LOADER>` prompt.

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

```
date
```

17. On node2, check the date by using the following command at the boot environment prompt:

```
show date
```

18. If necessary, set the date on node2:

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



Set the corresponding UTC date on node2.

19. On node2, check the time by using the following command at the boot environment prompt:

```
show time
```

20. If necessary, set the time on node2:

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



Set the corresponding UTC time on node2.

21. If necessary, set the partner system ID on node2:



This is the system ID of the node1 that you are upgrading to an AFF A900.

```
setenv partner-sysid node1_sysid
```

- a. Save the settings:

```
saveenv
```

22. On node2, at the LOADER prompt, verify the `partner-sysid` for node1:

```
printenv partner-sysid
```

For node2, the `partner-sysid` must be that of node1.

Stage 6. Boot node2 with the AFF A900 or FAS9500 controller and NVRAM modules

Boot node2 with AFF A900 or FAS9500 controller and NVRAM modules

Node2 with the AFF A900 or FAS9500 controller and NVRAM modules is now ready for upgrade. An AFF A700 to AFF A900 or FAS9000 to FAS9500 upgrade by swapping the controller and NVRAM modules involves moving only the console and management connections. This section provides the steps required to boot node2 with the AFF A900 or FAS9500 controller and NVRAM modules.

Steps

1. If NetApp Storage Encryption (NSE) is in use on this configuration, the `setenv bootarg.storageencryption.support` command must be set to `true`, and the `kmip.init.maxwait` variable must be set to `off` to avoid a boot loop after the node2 configuration is loaded:

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

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

2. Boot the node into the boot menu:

```
boot_ontap menu
```

3. Reassign the AFF A700 node2 disks to AFF A900 node2 by entering "22/7" and selecting the hidden option `boot_after_controller_replacement` when the node stops at the boot menu.

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

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

```
If [localhost:disk.encryptNoSupport:ALERT]: Detected FIPS-certified
encrypting drive and, or, [localhost:diskown.errorDuringIO:error]:
error 3 (disk failed) on disk errors occur, perform the following steps:
```



1. Halt the node at the LOADER prompt.
2. Check and reset the storage encryption bootargs mentioned in [Step 1](#).
3. At the loader prompt, boot up:

```
boot_ontap
```

You can use the following example as a reference:

Expand the console output example

```
LOADER-A> boot_ontap menu
.
.
<output truncated>
.
All rights reserved.
*****
*                                     *
* Press Ctrl-C for Boot Menu. *
*                                     *
*****
.
<output truncated>
.
Please choose one of the following:

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 22/7

(22/7)                               Print this secret List
(25/6)                               Force boot with multiple filesystem
disks missing.
(25/7)                               Boot w/ disk labels forced to clean.
(29/7)                               Bypass media errors.
(44/4a)                              Zero disks if needed and create new
flexible root volume.
(44/7)                               Assign all disks, Initialize all
disks as SPARE, write DDR labels
.
.
<output truncated>
.
.
(wipeconfig)                          Clean all configuration on boot
```

```

device
  (boot_after_controller_replacement) Boot after controller upgrade
  (boot_after_mcc_transition)         Boot after MCC transition
  (9a)                                Unpartition all disks and remove
their ownership information.
  (9b)                                Clean configuration and
initialize node with partitioned disks.
  (9c)                                Clean configuration and
initialize node with whole disks.
  (9d)                                Reboot the node.
  (9e)                                Return to main boot menu.

```

The boot device has changed. System configuration information could be lost. Use option (6) to restore the system configuration, or option (4) to initialize all disks and setup a new system. Normal Boot is prohibited.

Please choose one of the following:

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

This will replace all flash-based configuration with the last backup to disks. Are you sure you want to continue?: yes

.
.

<output truncated>

.
.

Controller Replacement: Provide name of the node you would like to replace:<nodename of the node being replaced>

Changing sysid of node nodel disks.

Fetches sanown old_owner_sysid = 536940063 and calculated old sys id

```
= 536940063
Partner sysid = 4294967295, owner sysid = 536940063
.
.
<output truncated>
.
.
varfs_backup_restore: restore using /mroot/etc/varfs.tgz
varfs_backup_restore: attempting to restore /var/kmip to the boot
device
varfs_backup_restore: failed to restore /var/kmip to the boot device
varfs_backup_restore: attempting to restore env file to the boot
device
varfs_backup_restore: successfully restored env file to the boot
device wrote key file "/tmp/rndc.key"
varfs_backup_restore: timeout waiting for login
varfs_backup_restore: Rebooting to load the new varfs
Terminated
<node reboots>

System rebooting...

.
.
Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...
Rebooting to load the restored env file...

.
System rebooting...

.
.
.
<output truncated>
.
.
.
.
WARNING: System ID mismatch. This usually occurs when replacing a
boot device or NVRAM cards!
Override system ID? {y|n} y
.
.
.
.
Login:
```



The system IDs shown in the preceding example are example IDs. The actual system IDs of the nodes that you are upgrading will be different.

Between entering node names at the prompt and the login prompt, the node reboots a few times to restore the environment variables, update firmware on the cards in the system, and for other ONTAP updates.

Verify the node2 installation

You must verify the node2 installation with the AFF A900 or FAS9500 controller and NVRAM modules. Because there is no change to physical ports, you are not required to map the physical ports from the old node2 to the replacement node2.

About this task

After you boot node1 with the AFF A900 or FAS9500 controller module, you must verify that it is installed correctly. You must wait for node2 to join quorum and then resume the controller replacement operation.

At this point in the procedure, the operation pauses while node2 joins quorum.

Steps

1. Verify that node2 has joined quorum:

```
cluster show -node node2 -fields health
```

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

2. Verify that node2 is part of the same cluster as node1 and that it is healthy:

```
cluster show
```

3. Switch to advanced privilege mode:

```
set advanced
```

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

```
system controller replace show
```

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

5. Resume the controller replacement operation:

```
system controller replace resume
```

6. The controller replacement operation pauses for intervention with the following message:

```

Cluster::*> system controller replace show
Node           Status           Error-Action
-----
Node2          Paused-for-intervention  Follow the instructions given
in
Node1          None              Step Details

```

Step Details:

To complete the Network Reachability task, the ONTAP network configuration must be manually adjusted to match the new physical network configuration of the hardware. This includes:

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



In this procedure, section *Re-creating VLANs, ifgrps, and broadcast domains* has been renamed *Restore network configuration on node2*.

7. With the controller replacement in a paused state, proceed to [Restore network configuration on node2](#).

Restore network configuration on node2

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

About this task

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

Steps

1. List all the physical ports that are on upgraded node2:

```
network port show -node node2
```

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

2. List the broadcast domains on the cluster:

```
broadcast-domain show
```

3. List network port reachability of all ports on node2:

```
network port reachability show -node node2
```

You should see output similar to the following example. The port and broadcast names vary.

```
Cluster::*> network port reachability show -node local
Node      Port      Expected Reachability      Reachability
Status
-----
Node2
      e0M      Default:Mgmt      no-reachability
      e10a      Default:Default-3      ok
      e10b      Default:Default-4      ok
      e11a      Cluster:Cluster      no-reachability
      e11b      Cluster:Cluster      no-reachability
      e11c      -      no-reachability
      e11d      -      no-reachability
      e2a      Default:Default-1      ok
      e2b      Default:Default-2      ok
      e9a      Default:Default      no-reachability
      e9b      Default:Default      no-reachability
      e9c      Default:Default      no-reachability
      e9d      Default:Default      no-reachability
13 entries were displayed.
```

In the preceding example, node2 has booted and joined quorum after controller replacement. It has several ports that have no reachability and are pending a reachability scan.

4. Repair the reachability for each of the ports on node2 with a reachability status other than `ok` by using the following command, in the following order:

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

- a. Physical ports
- b. VLAN ports

You should see output like the following example:

```
Cluster ::> reachability repair -node node2 -port e9d
```

```
Warning: Repairing port "node2:e9d" may cause it to move into a
different broadcast domain, which can cause LIFs to be re-homed away
from the port. Are you sure you want to continue? {y|n}:
```

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

Verify that all physical ports have their expected reachability:

```
network port reachability show
```

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

5. Verify port reachability:

```
network port reachability show
```

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

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

```
network port show
```

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

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

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

a. List any LIFs that are displaced:

```
displaced-interface show
```

b. Restore LIF home nodes and home ports:

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

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

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

Restore key-manager configuration on node2

If you are using NetApp Volume Encryption (NVE) to encrypt volumes on the system that you are upgrading, the encryption configuration must be synchronized to the new nodes. If you do not restore key-manager, when you relocate the node2 aggregates from the upgraded node1 to the upgraded node2 by using ARL, encrypted volumes will be taken offline.

About this task

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

Steps

1. Synchronize the encryption configuration for Onboard Key Manager (OKM) by using the following command at the cluster prompt:

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

2. Enter the cluster-wide passphrase for the OKM.

Move non-root aggregates and NAS data LIFs back to node2

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

About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. Moving SAN LIFs is not necessary for cluster or service health during the upgrade. SAN LIFs are not moved unless they need to be mapped to new ports. After you bring node2 online, you must verify that the LIFs are healthy and located on the appropriate ports.

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 back to node2, which is now running on the AFF A900 or the FAS9500 controller.

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

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

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

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

4. If necessary, restore and revert 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 node2:

```
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 7. Complete the upgrade

Confirm that the new controllers are set up correctly

To confirm the correct setup, you must verify that the HA pair is enabled. You must also verify that node1 and node2 can access each other's storage and that neither owns data LIFs belonging to other nodes on the cluster. In addition, you must verify that all data

aggregates are on their correct home nodes, and that the volumes for both nodes are online. If one of the new nodes has a unified target adapter, you must restore any port configurations and you might need to change the use of the adapter.

Steps

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

```
storage failover show
```

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

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

Node	Partner	Takeover Possible	State Description
node1	node2	true	Connected to node2
node2	node1	true	Connected to node1

3. Verify that node1 and node2 belong to the same cluster by using the following command and examining the output:

```
cluster show
```

4. Verify that node1 and node2 can access each other's storage by using the following command and examining the output:

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

5. Verify that neither node1 nor node2 owns data LIFs home-owned by other nodes in the cluster by using the following command and examining the output:

```
network interface show
```

If neither node1 or node2 owns data LIFs home-owned by other nodes in the cluster, revert the data LIFs to their home owner:

```
network interface revert
```

6. Verify that the aggregates are owned by their respective home nodes.

```
storage aggregate show -owner-name node1
```

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

7. Determine whether any volumes are offline:

```
volume show -node node1 -state offline
```

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

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

```
volume online -vserver vservice_name -volume volume_name
```

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

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

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

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

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

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

- Delete all expired licenses:

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

- Delete all unused licenses:

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

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

```
system license delete -serial-number node1_serial_number -package *  
system license delete -serial-number node2_serial_number -package *
```

The following output is displayed:

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

Enter *y* to remove all of the packages.

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

```
system license show
```

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

12. If NetApp Storage Encryption (NSE) was in use on the configuration and you set the `setenv`

`bootarg.storageencryption.support` command to `true` with the `<kmip.init.maxwait>` variable `off` (in *Boot node2 with the AFF A900 or the FAS9500 controller and NVRAM modules*, [Step 1](#)), you must reset the variable:

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

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

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

Refer to [References](#) to link to the *System Administration Reference* for information about the SPs and the *ONTAP 9 Commands: Manual Page Reference* for detailed information about the `system service-processor network modify` command.

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

After you finish

If Storage Encryption is enabled on node1 and node2, complete the section [Set up Storage Encryption on the new controller module](#). Otherwise, complete the section [Decommission the old system](#).

Set up Storage Encryption on the new controller module

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

About this task

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

Steps

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Set up NetApp Volume Encryption on the new controller module

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

About this task

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

Steps

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

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

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

- a. Add the key management server:

```
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:

For...	Use this command...
External key management (EKM)	<pre>security key-manager external restore</pre> <p>This command needs the OKM passphrase</p>
Onboard Key Manager	<pre>security key-manager onboard sync</pre>

After you finish

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

After you finish

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

Decommission the old system

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

Steps

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

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

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

4. Select **Go!**

A table displays cluster information, including the serial numbers.

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

Resume SnapMirror operations

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

Steps

1. Verify the SnapMirror status on the destination:

```
snapmirror show
```

2. Resume the SnapMirror relationship:

```
snapmirror resume -destination-vserver vserver_name
```

Troubleshoot

Troubleshoot

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

Refer to the table describing the different phases of the procedure in the section [Overview of the ARL upgrade](#). Information about the failures that can occur is listed by the phase of the procedure.

Aggregate relocation failures

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

Check for aggregate relocation failure

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

Steps

1. Enter the following command and examine the output:

```
storage aggregate relocation show
```

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

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

For detailed information about the `storage aggregate relocation start`, `override-vetoes`, and `override-destination-checks` options, refer to [References](#) to link to the *ONTAP 9 Commands: Manual Page Reference*.

Aggregates originally on node1 are owned by node2 after completion of the upgrade

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

About this task

Aggregates might fail to relocate correctly, that is, they have node2 as their home node instead of node1, under the following circumstances:

- During Stage 3, when aggregates are relocated from node2 to node1.

Some of the aggregates being relocated have node1 as their home node. For example, such an aggregate could be called `aggr_node_1`. If relocation of `aggr_node_1` fails during Stage 3, and relocation cannot be forced, then the aggregate is left behind on node2.

- After Stage 4, when node2 is replaced with the AFF A900 or the FAS9500 controller and NVRAM modules.

When node2 is replaced, `aggr_node_1` will come online with node1 as its home node instead of node2.

You can fix the incorrect ownership problem after Stage 6, after you have enabled storage failover by completing the following steps:

Steps

1. Get a list of aggregates:

```
storage aggregate show -nodes node2 -is-home true
```

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

2. Compare the output of Step 1 with the output you captured for node1 in the section [Prepare the nodes for upgrade](#) and note any aggregates that were not correctly relocated.
3. Relocate the aggregates left behind on node2:

```
storage aggregate relocation start -node node2 -aggr aggr_node_1 -destination node1
```

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

4. Verify that node1 is now the home owner of the aggregates:

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

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

Aggregates that do not have node1 as home owner can be relocated to node1 using the same relocation command in Step 3.

Reboots, panics, or power cycles

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

The solution to these problems depends on when they occur.

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

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

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

About this task

Takeover and giveback can proceed normally.

Steps

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

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

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

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

About this task

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

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

Steps

1. After node1 is booted up, all the non-root aggregates of node1 have moved back to node1. You must perform a manual aggregate relocation of the aggregates from node1 to node2:

```
storage aggregate relocation start -node node1 -destination node2 -aggregate  
-list * -ndocontroller-upgrade true
```
2. Continue with the node-pair upgrade procedure.

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

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

Steps

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

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

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

About this task

Node1 takes over all of node2's aggregates as well as any of its own aggregates that it had relocated to node2.

When node2 boots up, the aggregate relocation is completed automatically.

Steps

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

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

Node1 does not take over.

Steps

1. Bring up node2.

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

2. Continue the with rest of the node-pair upgrade procedure.

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

Node2 crashes during the first verification phase with HA pair disabled

Node1 does not take over following a node2 crash as the HA pair is already disabled.

Steps

1. Bring up node2.

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

2. Continue with the node-pair upgrade procedure.

Node1 crashes during the first verification phase with HA pair disabled

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

Steps

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

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

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

Node2 has relocated some or all of its aggregates from node1 to node1. Node1 serves data from aggregates that were relocated. The HA pair is disabled and hence there is no takeover.

About this task

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

Steps

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

Node1 crashes during the first resource-regain phase during aggregate relocation

If node1 crashes while node2 is relocating aggregates to node1, the task continues after node1 boots up.

About this task

Node2 continues to serve remaining aggregates, but aggregates that were already relocated to node1 encounter client outage while node1 is booting up.

Steps

1. Bring up node1.
2. Continue with the controller upgrade.

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

Node1 or node2 crashes during the post-check phase

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

Steps

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

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

Node1 crashes during the second resource-release phase

If node1 crashes while node2 is relocating aggregates, the task continues after node1 boots up.

About this task

Node2 continues to serve remaining aggregates but aggregates that were already relocated to node1 and node1's own aggregates encounter client outages while node1 is booting.

Steps

1. Bring up node1.
2. Continue with the controller upgrade procedure.

Node2 crashes during the second resource-release phase

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

About this task

Node1 continues to serve the aggregates that have been relocated, but the aggregates owned by node2 encounter client outages.

Steps

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

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

Node1 crashes during the second verification phase

If node1 crashes during this phase, takeover does not happen because the HA pair is already disabled.

About this task

There is a client outage for all aggregates until node1 reboots.

Steps

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

Node2 crashes during the second verification phase

If node2 crashes during this phase, takeover does not happen. Node1 serves data from the aggregates.

About this task

There is an outage for non-root aggregates that were already relocated until node2 reboots.

Steps

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

Issues that can arise in multiple stages of the procedure

Some issues can occur during different stages of the procedure.

Unexpected "storage failover show" command output

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

About this task

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

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

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

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

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

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

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

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

Steps

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

LIF migration failure

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

Steps

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

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

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

References

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

- [Reference content](#)

- [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.
High Availability management	Describes how to install and manage high-availability clustered configurations, including storage failover and takeover/giveback.
Logical storage management with the CLI	Describes how to efficiently manage your logical storage resources, using volumes, FlexClone volumes, files, and LUNs, FlexCache volumes, deduplication, compression, qtrees, and quotas.
MetroCluster Management and Disaster Recovery	Describes how to perform MetroCluster switchover and switchback operations, both in planned maintenance operations, or in the event of a disaster.
MetroCluster Upgrade and Expansion	Provides procedures for upgrading controller and storage models in the MetroCluster configuration, transitioning from a MetroCluster FC to a MetroCluster IP configuration, and expanding the MetroCluster configuration by adding additional nodes.
Network Management	Describes how to configure and manage physical and virtual network ports (VLANs and interface groups), LIFs, routing, and host-resolution services in clusters; optimize network traffic by load balancing; and monitor the cluster by using SNMP.
ONTAP 9.0 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.0 commands.
ONTAP 9.1 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.1 commands.
ONTAP 9.2 Commands: Manual Page Reference	Describes syntax and usage of supported ONTAP 9.2 commands.

Content	Description
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 AFF A700 to AFF A900 or FAS9000 to FAS9500 running ONTAP 9.10.1 or later	Describes the aggregate relocation procedures needed to non-disruptively upgrade an AFF A700 to an AFF A900 system or a FAS9000 to a FAS9500 system running ONTAP 9.10.1 or later by using "system controller replace" commands.
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 © 2022 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.