



Stage 3. Install and boot node3

Upgrade controllers

NetApp
March 06, 2026

Table of Contents

- Stage 3. Install and boot node3 1
 - Install and boot node3 1
 - Set the FC or UTA/UTA2 configuration on node3 6
 - Configure FC ports on node3 6
 - Check and configure UTA/UTA2 ports on node3 8
 - Reassign node1 disks to node3 10
 - Verify the node3 installation 16
 - Restore network configuration on node3 18
 - Restore key-manager configuration on node3 23
 - Move non-root aggregates and NAS data LIFs owned by node1 from node2 to node3 24

Stage 3. Install and boot node3

Install and boot node3

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

About this task

The relocation operation is paused at the beginning of this stage. This process is largely automated; the operation pauses to enable you to check its status. You must manually resume the operation. In addition, you must verify the SAN LIFs have successfully moved to node3.

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



- For an AFF A800 or AFF C800 controller upgrade, you must ensure that all drives in the chassis are firmly seated against the midplane before removing node1. For more information, see [Replace the AFF A800 or AFF C800 controller modules](#).
- If you are upgrading a system with storage disks, you need to complete this entire section and then go to the [Configure FC ports on node3](#) and [Check and configure UTA/UTA2 ports on node3](#) sections, entering commands at the cluster prompt.

Steps

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

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

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



If you're upgrading to a system with both nodes in the same chassis, install node4 and node3 in the chassis. If you don't install both nodes in the same chassis, when you boot node3, it behaves as if it were in a dual-chassis configuration, and when you boot node4, the interconnect between the nodes doesn't come up.

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

Cable the following connections using the *Installation and Setup Instructions* for the node3 platform, the appropriate disk shelf document, and the *HA pair management* documentation.

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

- Console (remote management port)
- Cluster ports

- Data ports
- Cluster and node management ports
- Storage
- SAN configurations: iSCSI Ethernet and FC switch ports



You might not need to move the interconnect card or the cluster interconnect cable connection from node1 to node3 because most platform models have a unique interconnect card model.

For the MetroCluster configuration, you need to move the FC-VI cable connections from node1 to node3. If the new host does not have an FC-VI card, you might need to move the FC-VI card.

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

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



When you boot node3, you might see the following warning message:

```
WARNING: The battery is unfit to retain data during a power outage. This is likely because the battery is discharged but could be due to other temporary conditions.
```

```
When the battery is ready, the boot process will complete and services will be engaged.
```

```
To override this delay, press 'c' followed by 'Enter'
```

5. If you see the warning message in [Step 4](#), take the following actions:
 - a. Check for any console messages that might indicate a problem other than a low NVRAM battery, and, if necessary, take any required corrective action.
 - b. Allow the battery to charge and the boot process to complete.



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



Refer to [Prepare for netboot](#).

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



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

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

7. Perform netboot on node3:

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

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



Do not interrupt the boot.

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

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

Disregard the following message:

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

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



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

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

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

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

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

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

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

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

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

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

```
ha-config show
```

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

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



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

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

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

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

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

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

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

14. Exit maintenance mode:

```
halt
```

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

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

```
date
```

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

```
show date
```

17. If necessary, set the date on node3:

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

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

```
show time
```

19. If necessary, set the time on node3:

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

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

```
setenv partner-sysid node2_sysid
```

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

- a. Save the settings:

```
saveenv
```

21. Verify the partner-sysid for node3:

```
printenv partner-sysid
```

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



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

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

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



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

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

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

23. Boot the node into boot menu:

```
boot_ontap menu
```

What's next?

- If you have a system with an FC or UTA/UTA2 configuration, [set and configure the FC or UTA/UTA2 ports on node3](#).
- If you don't have an FC or UTA/UTA2 configuration, [reassign node1 disks to node3, Step 1](#) so that node3 can recognize node1's disks.
- If you have a MetroCluster configuration, [set and configure the FC or UTA/UTA2 ports on node3](#) to detect the disks attached to the node.

Set the FC or UTA/UTA2 configuration on node3

If node3 has onboard FC ports, onboard unified target adapter (UTA/UTA2) ports, or a UTA/UTA2 card, you must configure the settings before completing the rest of the procedure.

About this task

You might need to complete the section [Configure FC ports on node3](#), the section [Check and configure UTA/UTA2 ports on node3](#), or both sections.



NetApp marketing materials might use the term UTA2 to refer to converged network adapter (CNA) adapters and ports. However, the CLI uses the term CNA.

If node3 does not have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card (for example, AFF and FAS systems introduced beginning with ONTAP 9.15.1), and you are upgrading a system with storage disks, you can skip to [Reassign node1 disks to node3](#).

Configure FC ports on node3

If node3 has FC ports, either onboard or on an add-on FC adapter, you must set port configurations on the

node before you bring it into service because the ports are not preconfigured when the systems are shipped. If you don't configure the ports, you might experience a disruption in service.

Before you begin

You must have the values of the FC port settings from node1 that you saved in the section [Prepare the nodes for upgrade](#).

About this task

You can skip this section if your system does not have FC configurations. If your system has onboard UTA/UTA2 ports or a UTA/UTA2 card, you configure them in [Check and configure UTA/UTA2 ports on node3](#).



Enter the commands in this section at the Maintenance mode shell prompt.

Steps

1. Compare the FC settings on node3 with the settings that you captured earlier from node1.
2. Take one of the following actions to modify the FC ports on node3, as needed:

In Maintenance mode (option 5 at boot menu):

- To program as target ports:

```
ucadmin modify -m fc -t target <adapter>
```

For example: `ucadmin modify -m fc -t target 2a`

- To program initiator ports:

```
ucadmin modify -m fc -t initiator <adapter>
```

For example: `ucadmin modify -m fc -t initiator 2b`

3. Verify the new settings by using the following command and examining the output:

```
ucadmin show
```

4. Halt the node:

```
halt
```

5. Boot the system from LOADER prompt:

```
boot_ontap menu
```

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

7. Select option 5 from the boot menu for maintenance mode.

8. Take one of the following actions:

If node 3...	Then...
Has a UTA/UTA2 card or UTA/UTA2 onboard ports	Go to Check and configure UTA/UTA2 ports on node3

If node 3...	Then...
Does not have a UTA/UTA2 card or UTA/UTA2 onboard ports	Skip <i>Check and configure UTA/UTA2 ports on node3</i> and go to Reassign node1 disks to node3 .

Check and configure UTA/UTA2 ports on node3

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

Before you begin

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

About this task

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



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

You can use the `ucadmin show` command to view or verify the current port configuration, as shown in the following example output:

```
*> ucadmin show
      Current  Current  Pending  Pending  Admin
Adapter Mode    Type    Mode    Type    Status
-----
0e     fc     target  -       initiator offline
0f     fc     target  -       initiator offline
0g     fc     target  -       initiator offline
0h     fc     target  -       initiator offline
1a     fc     target  -       -       online
1b     fc     target  -       -       online
6 entries were displayed.
```

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

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

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

request.



You must be in Maintenance mode to configure UTA/UTA2 ports. Enter the commands in this section at the Maintenance mode shell prompt.

Steps

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

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

2. Verify the UTA/UTA2 port settings:

```
ucadmin show
```

Examine the output and determine whether the UTA/UTA2 ports have the personality you want.

The output in the following example shows that the type of adapter "1b" is changing to initiator and that the mode of adapters "2a" and "2b" is changing to "cna". The CNA mode allows you to use the card as a network adapter.

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

3. Take one of the following actions:

If the UTA/UTA2 ports...	Then...
Do not have the personality that you want	Go to Step 4 .
Have the personality that you want	Skip Step 4 through Step 8 and go to Step 9 .

4. Take one of the following actions:

If you are configuring...	Then...
Ports on a UTA/UTA2 card	Go to Step 5
Onboard UTA/UTA2 ports	Skip Step 5 and go to Step 6 .

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

```
storage disable adapter <adapter_name>
```

Adapters in target mode are automatically offline in Maintenance mode.

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

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

- -m is the personality mode, fc or cna.
- -t is the FC4 type, target or initiator.



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

7. Place any target ports online by entering the following command once for each port:

```
storage enable adapter <adapter_name>
```

8. Cable the port.

9. Exit maintenance mode:

```
halt
```

10. Boot the node into boot menu:

```
boot_ontap menu
```

What's next?

- If you are upgrading to an AFF A800 system, go to [Reassign node1 disks to node3, Step 9](#).
- For all other system upgrades, go to [Reassign node1 disks to node3, Step 1](#).

Reassign node1 disks to node3

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

Steps

1. Verify that node1 has stopped at the boot menu. Reassign the disks of node1 to node3:

```
boot_after_controller_replacement
```

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

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

Expand the console output example

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

```
(9c) Clean configuration and
initialize node with whole disks.
(9d) Reboot the node.
(9e) Return to main boot menu.
The boot device has changed. System configuration information could
be lost. Use option (6) to restore the system configuration, or
option (4) to initialize all disks and setup a new system.
```

Normal Boot is prohibited.

Please choose one of the following:

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

Selection (1-11)? boot_after_controller_replacement

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

.

<output truncated>

.

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

Controller Replacement: Provide High Availability partner of node1:
<nodename of the partner of the node being replaced>

Changing sysid of node node1 disks.

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:

```



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

2. If the system goes into a reboot loop with the message `no disks found`, it indicates that the system has reset the FC or UTA/UTA2 ports back to the target mode and therefore is unable to see any disks. Select one of the following tasks to resolve this issue:

- Perform [Step 3](#) to [Step 8](#) on node3
- Go to section [Verify the node3 installation](#)

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

4. At the `LOADER` prompt, enter maintenance mode:

```
boot_ontap maint
```

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

```
ucadmin show
```

Change the ports back to initiator mode:

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

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

```
ucadmin show
```

7. Exit maintenance mode:

halt



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

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

8. At the LOADER prompt, boot up:

```
boot_ontap menu
```

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

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



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

- a. Return to the special boot menu:

```
LOADER> boot_ontap menu
```

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

Selection (1-11)? 10
```

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

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

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

- d. At the prompt, enter the key-manager passphrase.
- e. Enter the backup data when prompted.



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

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



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

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



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

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

- a. Enter maintenance mode:

```
boot_ontap maint
```

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

```
aggr status -r
```

- c. Check the status of the node1 aggregate:

```
aggr status
```

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

```
aggr_online root_aggr_from_node1
```

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

```
aggr offline root_aggr_on_node3
```

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

```
aggr options aggr_from_node1 root
```

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

```
aggr status
```



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

The following shows an example of the command output:

```

-----
Aggr                State      Status      Options
aggr0_nst_fas8080_15 online    raid_dp, aggr fast zeroed
                    64-bit
aggr0                offline   raid_dp, aggr diskroot
                    fast zeroed
                    64-bit
-----

```

Verify the node3 installation

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

About this task

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

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

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

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

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

Steps

1. Verify that node3 has joined quorum:

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

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

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

```
cluster show
```

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

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

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

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

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

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

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

- Switch to advanced privilege mode:

```
set advanced
```

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

```
system controller replace show

system controller replace show-details
```

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

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

Check the MetroCluster node status:

```
metrocluster node show
```

8. Resume the controller replacement operation:

```
system controller replace resume
```

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

```
Cluster::*> system controller replace show
Node          Status          Error-Action
-----
Node1(now node3) Paused-for-intervention  Follow the instructions
given in
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, the section *Re-creating VLANs, ifgrps, and broadcast domains* has been renamed *Restore network configuration on node3*.

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

Restore network configuration on node3

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

About this task

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



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

Steps

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

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

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

2. List the broadcast domains on the cluster:

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

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

```
network port reachability show
```

You should see output like the following example:

```

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

```

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

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

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

You should see output like the following example:

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

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

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

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

Verify that all physical ports have their expected reachability:

```
network port reachability show
```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

a. List displaced VLANs:

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

Output like the following should display:

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

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

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

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

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

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

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

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

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

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

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

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

10. Verify port reachability:

```
network port reachability show
```

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

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

```
network port show
```

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

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

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

- a. List any LIFs that are displaced:

```
displaced-interface show
```

- b. Restore LIF home nodes and home ports:

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

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

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

Restore key-manager configuration on node3

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

About this task

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

Steps

1. Run the following command from node3:

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

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

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

Example

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

node	vserver	key-server	key-id
restored			
-----	-----	-----	-----
node3	svm1	""	0000000000000000020000000000a008a81976
true			2190178f9350e071fbb90f0000000000000000

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

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

About this task

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



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

Steps

1. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

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

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

2. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following checks:

- Cluster health check
- Cluster LIF status check

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

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

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

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

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

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

Restore and revert displaced LIFs

- a. List any displaced LIFs:

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

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

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

Manually migrate and modify LIFs

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

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

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

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

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

```
system controller replace resume
```

The system performs the following post-checks:

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

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.