



Stage 5. Install and boot node4

Upgrade controllers

NetApp

February 10, 2026

Table of Contents

Stage 5. Install and boot node4	1
Install and boot node4	1
Set the FC or UTA/UTA2 configuration on node4	6
Configure FC ports on node4	6
Check and configure UTA/UTA2 ports on node4	7
Reassign node2 disks to node4	10
Map ports from node2 to node4	15
Join the quorum when a node has a different set of network ports	19
Verify the node4 installation	20
Re-create VLANs, interface groups, and broadcast domains on node4	21
Restore key-manager configuration on node4	21
Move non-root aggregates and NAS data LIFs owned by node2 from node3 to node4	22

Stage 5. Install and boot node4

Install and boot node4

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

About this task

The relocation operation is paused at the beginning of this stage. This process is mostly automated; the operation pauses to enable you to check its status. You must manually resume the operation. In addition, you must verify that the NAS data LIFs have successfully moved to node4.

You need to netboot node4 if the ONTAP version on node4 is different to the ONTAP version on node2. After you install node4, boot it from the ONTAP 9 image stored on the web server. You can then download the correct files to the boot media device for subsequent system boots by following the instructions in [Prepare for netboot](#).

- For an AFF A800 or AFF C800 controller upgrade, you must ensure that all drives in the chassis are firmly seated against the midplane before removing node2. For more information, see [Replace the AFF A800 or AFF C800 controller modules](#).
- If you are upgrading a system with storage disks, you must complete this entire section and then proceed to [Set the FC or UTA/UTA2 configuration on node4](#), entering commands at the cluster prompt.

Steps

1. Make sure that node4 has sufficient rack space.

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

2. Install node4 in the rack, following the instructions in the *Installation and Setup Instructions* for the node model.
3. Cable node4, moving the connections from node2 to node4.

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

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

- Console (remote management port)
- Cluster ports
- Data ports
- Cluster and node management ports
- Storage
- SAN configurations: iSCSI Ethernet and FC switch ports



You might not need to move the interconnect card/FC-VI card or interconnect/FC-VI cable connection from node2 to node4 because most platform models have unique interconnect card models.

For the MetroCluster configuration, you must move the FC-VI cable connections from node2 to node4. If the new host does not have an FC-VI card, you might need to move the FC-VI card.

4. Turn on the power to node4, and then interrupt the boot process by pressing Ctrl-C at the console terminal to access the boot environment prompt.



When you boot node4, you might see the following warning message:

```
WARNING: The battery is unfit to retain data during a power outage. This
is likely
because the battery is discharged but could be due to other
temporary
conditions.
When the battery is ready, the boot process will complete
and services will be engaged. To override this delay, press 'c'
followed
by 'Enter'
```

5. If you see the warning message in Step 4, take the following actions:

- a. Check for any console messages that might indicate a problem other than a low NVRAM battery, and, if necessary, take any required corrective action.
- b. Allow the battery to charge and the boot process to complete.



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



Refer to [Prepare for netboot](#).

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



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

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

If Dynamic Host Configuration Protocol (DHCP) is...	Then...
Not running	<p>Manually configure the connection by entering the following command at the boot environment prompt:</p> <pre>ifconfig e0M -addr=filer_addr -mask=netmask - gw=gateway -dns=dns_addr -domain=dns_domain</pre> <p><i>filer_addr</i> is the IP address of the storage system (mandatory). <i>netmask</i> is the network mask of the storage system (mandatory). <i>gateway</i> is the gateway for the storage system (mandatory). <i>dns_addr</i> is the IP address of a name server on your network (optional). <i>dns_domain</i> is the Domain Name Service (DNS) domain name. If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL; you need only the server's host name.</p> <p>NOTE: Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p>

7. Perform netboot on node4:

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

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



Do not interrupt the boot.

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

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

Disregard the following message:

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

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



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

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

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

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

- Enter **n** to skip the backup recovery when you see the following prompt:

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

- Reboot by entering **y** when you see the following prompt:

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

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

- Select maintenance mode 5 from the boot menu and enter **y** when you are prompted to continue with the boot.
- 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.

- 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
```

- Exit maintenance mode:

```
halt
```

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

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

```
date
```

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

```
show date
```

17. If necessary, set the date on node4:

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

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

```
show time
```

19. If necessary, set the time on node4:

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

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

```
setenv partner-sysid node3_sysid
```

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

Save the settings:

```
saveenv
```

21. Verify the partner-sysid for node4:

```
printenv partner-sysid
```

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



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

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

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



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

b. Contact NetApp Support for assistance with restoring the onboard key management information.

23. Boot the node into boot menu:

```
boot_ontap menu
```

What's next?

- If you have a system with an FC or UTA/UTA2 configuration, [set the FC or UTA/UTA2 configuration on node4](#).
- If you don't have an FC or UTA/UTA2 configuration, [reassign node2 disks to node4, Step 1](#) so that node4 can recognize node2's disks.
- If you have a MetroCluster configuration, [set the FC or UTA/UTA2 configuration on node4](#) to detect the disks attached to the node.

Set the FC or UTA/UTA2 configuration on node4

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

About this task

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

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



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

Configure FC ports on node4

If node4 has FC ports, either onboard or on an add-on FC adapter, you must set port configurations on the node before you bring it into service because the ports are not preconfigured when the systems are shipped. If you don't configure the ports as required, you might experience a disruption in service.

Before you begin

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

About this task

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



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

Steps

1. Display information about all FC and converged network adapters on the system:

```
system node hardware unified-connect show
```

2. Compare the FC settings on node4 with the settings that you captured earlier from node1.

3. Modify the FC ports on node4 as needed:

- To program as target ports:

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

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

- To program initiator ports:

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

`-t` is the FC4 type: target or initiator.

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

4. Halt the node:

```
halt
```

5. Boot the system from LOADER prompt:

```
boot_ontap menu
```

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

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

8. Take one of the following actions:

- Go to [Check and configure UTA/UTA2 ports on node4](#) if node4 has a UTA/UTA2 card or UTA/UTA2 onboard ports.
- If node4 doesn't have a UTA/UTA2 card or UTA/UTA2 onboard ports, skip [Check and configure UTA/UTA2 ports on node4](#) and go to [Reassign node2 disks to node4](#).

Check and configure UTA/UTA2 ports on node4

If node4 has onboard UTA/UTA2 ports or a UTA/UTA2A card, you must check the configuration of the ports and configure them, depending on how you want to use the upgraded system.

Before you begin

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

About this task

UTA/UTA2 ports can be configured into native FC mode or UTA/UTA2A mode. FC mode supports FC initiator

and FC target; UTA/UTA2 mode allows concurrent NIC and FCoE traffic to share the same 10GbE SFP+ interface and supports FC target.



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

UTA/UTA2 ports might be on an adapter or on the controller with the following configurations:

- UTA/UTA2 cards ordered at the same time as the controller are configured before shipment to have the personality you requested.
- UTA/UTA2 cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard UTA/UTA2 ports on new controllers are configured (before shipment) to have the personality you requested.

However, you should check the configuration of the UTA/UTA2 ports on node4 and change it, if necessary.



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

Steps

1. Check how the ports are currently configured on node4:

```
system node hardware unified-connect show
```

The system displays output similar to the following example:

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

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

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

3. Verify the settings:

```
ucadmin show
```

Examine the output of the `ucadmin show` command and determine whether the UTA/UTA2 ports have the personality you want.

The output in the following examples shows that the FC4 type of adapter "1b" is changing to initiator and that the mode of adapters "2a" and "2b" is changing to cna:

```
*> ucadmin show
Node Adapter Current Mode Current Type Pending Mode Pending Type
Admin Status
---- ----- ----- -----
-----
f-a 1a fc initiator -
online
f-a 1b fc target -
online
f-a 2a fc target cna -
online
f-a 2b fc target cna -
online
4 entries were displayed.
*>
```

4. Take one of the following actions:

If the CNA ports...	Then...
Do not have the personality that you want	Go to Step 5 .
Have the personality that you want	Skip Step 5 through Step 9 and go to Step 10 .

5. Take one of the following actions:

If you are configuring...	Then...
Ports on a UTA/UTA2 card	Go to Step 6
Onboard UTA/UTA2 ports	Skip Step 6 and go to Step 7 .

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

```
storage disable adapter adapter_name
```

Adapters in target mode are automatically offline in Maintenance mode.

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

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

- *-m* is the personality mode, FC or 10GbE UTA.
- *-t* is the FC4 type, target or initiator.



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

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

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

9. Cable the port.

10. Exit Maintenance mode:

```
halt
```

11. Boot the node into boot menu:

```
boot_ontap menu
```

What's next?

- If you are upgrading to an AFF A800 system, go to [Reassign node2 disks to node4, Step 9](#).
- For all other system upgrades, go to [Reassign node2 disks to node4, Step 1](#).

Reassign node2 disks to node4

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

Steps

1. Verify that node2 has stopped at the boot menu and reassign the disks of node2 to node4:

```
boot_after_controller_replacement
```

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

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

Expand the console output example

```
LOADER-A> boot_ontap menu ...
*****
*                                     *
* Press Ctrl-C for Boot Menu. *
*                                     *
*****
.

.

Please choose one of the following:

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

.

.

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

Please choose one of the following:

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
Selection (1-9)? boot_after_controller_replacement
```

```

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

.
.

Controller Replacement: Provide name of the node you would like to
replace: <name of the node being replaced>
Controller Replacement: Provide High Availability partner of node1:
<nodename of the partner of the node being replaced>
Changing sysid of node <node being replaced> disks.
Fetched sanown old_owner_sysid = 536953334 and calculated old sys id
= 536953334
Partner sysid = 4294967295, owner sysid = 536953334

.
.

.
.

Terminated
<node reboots>

.
.

System rebooting...

.
.

Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...

.
.

System rebooting...

.
.

.
.

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

```

2. If the system goes into a reboot loop with the message no disks found, this is because it has reset the ports back to the target mode and therefore is unable to see any disks. Perform [Step 3](#) through [Step 8](#) on node4 to resolve this issue.
3. Press Ctrl-C during AUTOBOOT to stop the node at the LOADER> prompt.
4. At the LOADER prompt, enter maintenance mode:

```
boot_ontap maint
```

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

```
ucadmin show
```

Change the ports back to initiator mode:

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

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

```
ucadmin show
```

7. Exit maintenance mode:

```
halt
```

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



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

8. At the LOADER prompt, boot up:

```
boot_ontap menu
```

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

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

a. Return to the special boot menu:

```
LOADER> boot_ontap menu
```

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

```
Selection (1-11)? 10
```

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

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

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

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

e. Enter the backup data when prompted.



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

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



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

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



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

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

a. Enter maintenance mode:

```
boot_ontap maint
```

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

```
aggr status -r
```

c. Check the status of the node2 aggregate:

```
aggr status
```

d. If necessary, bring the node2 aggregate online:

```
aggr_online root_aggr_from_node2
```

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

```
aggr_offline root_aggr_on_node4
```

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

```
aggr_options aggr_from_node2 root
```

Map ports from node2 to node4

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

About this task

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

The software configuration of node4 must match the physical connectivity of node4 and IP connectivity must be restored before you continue with the upgrade.

Port settings might vary, depending on the model of the nodes. You must make the original node's port and LIF configuration compatible with what you plan the new node's configuration to be. This is because the new node replays the same configuration when it boots, meaning when you boot node4 that Data ONTAP will try to host LIFs on the same ports that were used on node2.

Therefore, if the physical ports on node2 do not map directly to the physical ports on node4, then software configuration changes will be required to restore cluster, management, and network connectivity after the boot. In addition, if the cluster ports on node2 do not directly map to the cluster ports on node4, node4 might not automatically rejoin the quorum when it is rebooted until a software configuration change is made to host the cluster LIFs on the correct physical ports.

Steps

1. Record all the node2 cabling information for node2, the ports, broadcast domains, and IPspaces, in the table:

LIF	Node2 ports	Node2 IPspaces	Node2 broadcast domains	Node4 ports	Node4 IPspaces	Node4 broadcast domains
Cluster 1						
Cluster 2						
Cluster 3						
Cluster 4						
Node management						
Cluster management						
Data 1						
Data 2						
Data 3						
Data 4						
SAN						
Intercluster port						

2. Record all the cabling information for node4, the ports, broadcast domains, and IPspaces, in the table.

3. Follow these steps to verify if the setup is a two-node switchless cluster:

a. Set the privilege level to advanced:

```
cluster::> set -privilege advanced
```

b. Verify if the setup is a two-node switchless cluster:

```
cluster::> network options switchless-cluster show
```

```
cluster::*> network options switchless-cluster show
Enable Switchless Cluster: false/true
```

The value of this command must match the physical state of the system.

c. Return to the administration privilege level:

```
cluster::*> set -privilege admin
cluster::>
```

4. Follow these steps to place node4 into quorum:

a. Boot node4. See [Install and boot node4](#) to boot the node if you have not already done so.

b. Verify that the new cluster ports are in the Cluster broadcast domain:

```
network port show -node node -port port -fields broadcast-domain
```

The following example shows that port "e0a" is in the Cluster domain on node4:

```
cluster::> network port show -node node4 -port e0a -fields broadcast-
domain
node      port broadcast-domain
-----
node4    e0a  Cluster
```

c. If the cluster ports are not in the Cluster broadcast-domain, add them with the following command:

```
broadcast-domain add-ports -ipspace Cluster -broadcast-domain Cluster -ports
node:port
```

d. Add the correct ports to the Cluster broadcast domain:

```
network port modify -node -port -ipspace Cluster -mtu 9000
```

This example adds Cluster port "e1b" on node4:

```
network port modify -node node4 -port e1b -ipspace Cluster -mtu 9000
```

e. Migrate the cluster LIFs to the new ports, once for each LIF:

```
network interface migrate -vserver Cluster -lif lif_name -source-node node4  
destination-node node4 -destination-port port_name
```

f. Modify the home port of the cluster LIFs:

```
network interface modify -vserver Cluster -lif lif_name -home-port port_name
```

g. Remove the old ports from the Cluster broadcast domain:

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

This command removes port "e0d" on node4:

```
network port broadcast-domain remove-ports -ipspace Cluster -broadcast  
-domain Cluster -ports node4:e0d
```

h. Verify that node4 has rejoined quorum:

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

5. Adjust the broadcast domains hosting your cluster LIFs and node-management/clustermanagement LIFs. Confirm that each broadcast domain contains the correct ports. A port cannot be moved between broadcast domains if it is hosting or is home to a LIF so you may need to migrate and modify the LIFs as shown in the following steps:

a. Display the home port of a LIF:

```
network interface show -fields home-node,home-port
```

b. Display the broadcast domain containing this port:

```
network port broadcast-domain show -ports node_name:port_name
```

c. Add or remove ports from broadcast domains:

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

d. Modify a LIF's home port:

```
network interface modify -vserver vserver -lif lif_name -home-port port_name
```

6. Adjust the intercluster broadcast domains and migrate the intercluster LIFs, if necessary, using the same commands shown in [Step 5](#).

7. Adjust any other broadcast domains and migrate the data LIFs, if necessary, using the same commands shown in [Step 5](#).

8. If there were any ports on node2 that no longer exist on node4, follow these steps to delete them:

a. Access the advanced privilege level on either node:

```
set -privilege advanced
```

b. To delete the ports:

```
network port delete -node node_name -port port_name
```

c. Return to the admin level:

```
set -privilege admin
```

9. Adjust all the LIF failover groups:

```
network interface modify -failover-group failover_group -failover-policy failover_policy
```

The following command sets the failover policy to broadcast-domain-wide and uses the ports in failover group fg1 as failover targets for LIF data1 on node4:

```
network interface modify -vserver node4 -lif data1 failover-policy broadcast-domainwide -failover-group fg1
```

Refer to [References](#) to link to *Network Management* or the *ONTAP 9 Commands: Manual Page Reference* and see *Configuring failover settings on a LIF* for more information.

10. Verify the changes on node4:

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

11. Each cluster LIF must be listening on port 7700. Verify that the cluster LIFs are listening on port 7700:

```
::> network connections listening show -vserver Cluster
```

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

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

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

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

Repeat step 11 to verify that the cluster LIF is now listening on port 7700.

Join the quorum when a node has a different set of network ports

The node with the new controller boots and attempts to join the cluster automatically at first; however, if the new node has a different set of network ports, you must perform the following steps to confirm that the node successfully joins the quorum.

About this task

You can use these instructions for any relevant node. Node3 is used throughout the following sample.

Steps

1. Verify that the new cluster ports are in the Cluster broadcast domain by entering the following command and checking the output:

```
network port show -node node -port port -fields broadcast-domain
```

The following example shows that port "e1a" is in the Cluster domain on node3:

```
cluster::> network port show -node node3 -port e1a -fields broadcast-
domain
node      port      broadcast-domain
-----  -----
node3    e1a      Cluster
```

2. Add the correct ports to the Cluster broadcast domain by entering the following command and checking the output:

```
network port modify -node -port -ipspace Cluster -mtu 9000
```

This example adds Cluster port "e1b" on node3:

```
network port modify -node node3 -port e1b -ipspace Cluster -mtu 9000
```

3. Migrate the cluster LIFs to the new ports, once for each LIF, using the following command:

```
network interface migrate -vserver Cluster -lif lif_name -source-node node3
destination-node node3 -destination-port port_name
```

4. Modify the home port of the cluster LIFs as follows:

```
network interface modify -vserver Cluster -lif lif_name -home-port port_name
```

5. If the cluster ports are not in the Cluster broadcast-domain, add them with the following command:

```
network port broadcast-domain add-ports -ipspace Cluster -broadcastdomain
```

```
Cluster ports node:port
```

6. Remove the old ports from the Cluster broadcast domain. You can use for any relevant node. The following command removes port "e0d" on node3:

```
network port broadcast-domain remove-ports network port broadcast-domain
remove-ports ipspace Cluster -broadcast-domain Cluster -ports node3:e0d
```

7. Verify the node has rejoined quorum as follows:

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

8. Adjust the broadcast domains hosting your cluster LIFs and node-management/cluster management LIFs. Confirm that each broadcast domain contains the correct ports. A port cannot be moved between broadcast domains if it is hosting or is home to a LIF, so you might need to migrate and modify the LIFs as follows:

- a. Display the home port of a LIF:

```
network interface show -fields home-node,home-port
```

- b. Display the broadcast domain containing this port:

```
network port broadcast-domain show -ports node_name:port_name
```

- c. Add or remove ports from broadcast domains:

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

- d. Modify a home port of a LIF:

```
network interface modify -vserver vserver-name -lif lif_name -home-port
port_name
```

Adjust the intercluster broadcast domains and migrate the intercluster LIFs, if necessary. The data LIFs remain unchanged.

Verify the node4 installation

After you install and boot node4, you must verify that it is installed correctly, that it is part of the cluster, and that it can communicate with node3.

About this task

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

Steps

1. Verify that node4 has joined quorum:

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

2. Verify that node4 is part of the same cluster as node3 and healthy by entering the following command:

```
cluster show
```

3. Check the status of the operation and verify that the configuration information for node4 is the same as node2:

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

If the configuration is different for node4, a system disruption might occur later in the procedure.

4. Check that the replaced controller is configured correctly for MetroCluster configuration and not in switch-over mode.



At this stage MetroCluster configuration will not be in a normal state and you might have errors to resolve. See [Verify the health of the MetroCluster configuration](#).

Re-create VLANs, interface groups, and broadcast domains on node4

After you confirm that node4 is in quorum and can communicate with node3, you must re-create node2's VLANs, interface groups, and broadcast domains on node4. You must also add the node3 ports to the newly re-created broadcast domains.

About this task

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

Steps

1. Re-create the VLANs on node4 using the node2 information recorded in the [Relocate non-root aggregates and NAS data LIFs from node2 to node3](#) section:

```
network port vlan create -node node4 -vlan vlan-names
```

2. Re-create the interface groups on node4 using the node2 information recorded in the [Relocate non-root aggregates and NAS data LIFs from node2 to node3](#) section:

```
network port ifgrp create -node node4 -ifgrp port_ifgrp_names-distr-func
```

3. Re-create the broadcast domains on node4 using the node2 information recorded in the [Relocate non-root aggregates and NAS data LIFs from node2 to node3](#) section:

```
network port broadcast-domain create -ipspace Default -broadcast-domain broadcast_domain_names -mtu mtu_size -ports node_name:port_name,node_name:port_name
```

4. Add the node4 ports to the newly re-created broadcast domains:

```
network port broadcast-domain add-ports -broadcast-domain broadcast_domain_names -ports node_name:port_name,node_name:port_name
```

Restore key-manager configuration on node4

If you are using NetApp Aggregate Encryption (NAE) or NetApp Volume Encryption (NVE) to encrypt volumes on the system you are upgrading, the encryption configuration must be synchronized to the new nodes. If you

do not restore key-manager, when you relocate the node2 aggregates from node3 to node4 by using ARL, encrypted volumes will be taken offline.

Steps

1. To synchronize encryption configuration for Onboard Key Manager, run the following command at the cluster prompt:

For this ONTAP version...	Use this command...
ONTAP 9.6 or 9.7	<code>security key-manager onboard sync</code>
ONTAP 9.5	<code>security key-manager setup -node node_name</code>

2. Enter the cluster-wide passphrase for the Onboard Key Manager.

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

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

About this task

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

Steps

1. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

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

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

2. Manually verify that the network and all VLANs, interface groups, and broadcast domains have been configured correctly.
3. Resume the relocation operation:

```
system controller replace resume
```

To complete the "Network Reachability" phase, ONTAP network configuration must be manually adjusted to match the new physical network configuration of the hardware. This includes assigning network ports to the correct broadcast domains, creating any required ifgrps and VLANs, and modifying the home-port parameter of network interfaces to the appropriate ports. Refer to the "Using aggregate relocation to upgrade controller hardware on a pair of nodes running ONTAP 9.x" documentation, Stages 3 and 5. Have all of these steps been manually completed? [y/n]

4. Enter y to continue.
5. The system performs the following checks:

- Cluster health check
- Cluster LIF status check

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

The system pauses once the resource relocation is complete.

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

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

7. Manually verify that the non-root aggregates and NAS data LIFs have been successfully relocated to node4.

If any aggregates fail to relocate or are vetoed, you must take manually relocate the aggregates, or override either the vetoes or destination checks, if necessary. See the section [Relocate failed or vetoed aggregates](#) for more information.

8. Confirm that the SAN LIFs are on the correct ports on node4 by completing the following substeps:

- a. Enter the following command and examine its output:

```
network interface show -data-protocol iscsi|fcp -home-node node4
```

The system returns output similar to the following example:

```

cluster::> net int show -data-protocol iscsi|fcp -home-node node3
      Logical      Status      Network          Current Current Is
  Vserver  Interface  Admin/Oper  Address/Mask      Node      Port      Home
  -----
vs0
  a0a        up/down    10.63.0.53/24    node3    a0a    true
  data1      up/up      10.63.0.50/18    node3    e0c    true
  rads1      up/up      10.63.0.51/18    node3    e1a    true
  rads2      up/down    10.63.0.52/24    node3    e1b    true
vs1
  lif1      up/up      172.17.176.120/24  node3    e0c    true
  lif2      up/up      172.17.176.121/24  node3    e1a    true

```

b. If node4 has any SAN LIFs or groups of SAN LIFs that are on a port that did not exist on node2 or that need to be mapped to a different port, move them to an appropriate port on node4 by completing the following substeps:

i. Set the LIF status to down by entering the following command:

```
network interface modify -vserver vserver_name -lif lif_name -status
-admin down
```

ii. Remove the LIF from the port set:

```
portset remove -vserver vserver_name -portset portset_name -port-name
port_name
```

iii. Enter one of the following commands:

- Move a single LIF by entering the following command:

```
network interface modify -vserver vserver_name -lif lif_name -home
-port new_home_port
```

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

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

- Add the LIFs back to the port set:

```
portset add -vserver vserver_name -portset portset_name -port-name
port_name
```



You must confirm that you move SAN LIFs to a port that has the same link speed as the original port.

c. Modify the status of all LIFs to up so the LIFs can accept and send traffic on the node by entering the following command:

```
network interface modify -home-port port_name -home-node node4 -lif data  
-statusadmin up
```

d. Enter the following command and examine its output to verify that LIFs have been moved to the correct ports and that the LIFs have the status of *up* by entering the following command on either node and examining the output:

```
network interface show -home-node <node4> -role data
```

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

```
network interface modify -vserver vserver_name -lif lif_name -status-admin  
up
```

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

```
system controller replace resume
```

The system performs the following post-checks:

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

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—with 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.