



Stage 5. Install and boot node4

Upgrade controllers

NetApp
March 06, 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 9
 - Configure FC ports on node4 10
 - Check and configure UTA/UTA2 ports on node4 11
 - Map ports from node2 to node4 13
 - Verify the node4 installation 18
 - Move NAS data LIFs owned by node2 from node3 to node4 and verify SAN LIFs on node4 19
 - Worksheet: Information to record before moving NAS data LIFs to node4 26
 - Relocate node2 non-root aggregates from node3 to node4 27

Stage 5. Install and boot node4

Install and boot node4

You must install node4 in the rack, transfer node2 connections to node4, and boot node4. You must also reassign any node2 spares, any disks belonging to root, and any non-root aggregates that were not relocated to node3 earlier.

About this task

You must 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](#)

However, you don't need to netboot node4 if the ONTAP version on node4 is the same or later than the ONTAP version on node2.



- 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're upgrading a system with storage disks, you must complete this entire section and then proceed to the section [Set the FC or UTA/UTA2 configuration on node4](#), entering commands at the cluster prompt.

Steps

1. Take one of the following actions:

If node4 will be in ...	Then...
A chassis separate from node3	Go to Step 2 .
The same chassis with node3	Skip Steps 2 and 3 and go to Step 4 .

2. Make sure that node4 has sufficient rack space.

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

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

The following references help you make proper cable connections. Go to [References](#) to link to them.

- *Installation and Setup Instructions* for the node4 platform
- The appropriate disk shelf procedure
- The *HA pair management* documentation

Cable the following connections:

- Console (remote management port)

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



You do 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.

5. Take one of the following actions:

If node4 is in...	Then...
The same chassis as node3	Go to Step 8 .
A chassis separate from node3	Go to Step 6 .

6. Turn on the power to node4, and then interrupt the boot by pressing Ctrl-C to access the boot environment prompt.



When you boot node4, you might see the following 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'
```

7. If you see the warning message in Step 6, take the following actions:
- Check for any console messages that might indicate a problem other than a low NVRAM battery and, if necessary, take any required corrective action.
 - Allow the battery to charge and the boot process to finish.



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

8. At the Maintenance mode prompt, enter the following command:

```
halt
```

The system stops at the boot environment prompt.

9. Configure node4 for ONTAP:

```
set-defaults
```

10. 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.

11. If the version of ONTAP installed on node4 is the same or later than the version of ONTAP 9 installed on node2, enter the following command:

```
boot_ontap menu
```

12. Take one of the following actions:

If the system you are upgrading...	Then...
Does not have the correct or current ONTAP version on node4	Go to Step 13 .
Has the correct or current version of ONTAP on node4	Go to Step 18 .

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



You must use the management port and IP address as the netboot connection. Do not use a data LIF IP address 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=<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. 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> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details. </div>

14. Perform netboot on node4:

For...	Then...
FAS/AFF8000 series systems	<pre>netboot http://<web_server_ip/path_to_webaccessible_directory> /netboot/kernel</pre>
All other systems	<pre>netboot http://<web_server_ip/path_to_webaccessible_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 [Step 1](#) in the section *Prepare for netboot*.

 Do not interrupt the boot.

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

This menu option downloads and installs the new Data 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 Data 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 releases of ONTAP. The netboot procedure combined with option (7) `Install new software` wipes the boot media and places the same ONTAP version on both image partitions.

16. 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
```

17. Complete the following substeps:

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

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

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

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

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

18. Select maintenance mode 5 from the boot menu and enter `y` when you are prompted to continue with the boot.
19. Before continuing, go to [Set the FC or UTA/UTA2 configuration on node4](#) to make any necessary changes to the FC or UTA/UTA2 ports on the node. Make the changes recommended in those sections, reboot the node, and go into Maintenance mode.
20. Enter the following command and examine the output to find the system ID of node4:

```
disk show -a
```

The system displays the system ID of the node and information about its disks, as shown in the following example:

```

*> disk show -a
Local System ID: 536881109
DISK          OWNER                                POOL  SERIAL NUMBER  HOME
-----
0b.02.23     nst-fas2520-2 (536880939)  Pool0 KPG2RK6F      nst-
fas2520-2 (536880939)
0b.02.13     nst-fas2520-2 (536880939)  Pool0 KPG3DE4F      nst-
fas2520-2 (536880939)
0b.01.13     nst-fas2520-2 (536880939)  Pool0 PPG4KLAA      nst-
fas2520-2 (536880939)
.....
0a.00.0      (536881109)                Pool0 YFKSX6JG
(536881109)
.....

```

21. Reassign node2's spares, disks belonging to the root, and any non-root aggregates that were not relocated to node3 earlier in section [Relocate non-root aggregates from node2 to node3](#):



If you have shared disks, hybrid aggregates, or both on your system, you must use the correct `disk reassign` command from the following table.

Disk type...	Run the command...
With shared disks	<pre> disk reassign -s node2_sysid -d node4_sysid -p node3_sysid </pre>
Without shared	<pre> disks disk reassign -s node2_sysid -d node4_sysid </pre>

For the `<node2_sysid>` value, use the information captured in [Step 10](#) of the *Record node2 information* section. For `node4_sysid`, use the information captured in [Step 23](#).



The `-p` option is only required in maintenance mode when shared disks are present.

The `disk reassign` command will reassign only those disks for which `node2_sysid` is the current owner.

The system displays the following message:

```
Partner node must not be in Takeover mode during disk reassignment from
maintenance mode.
Serious problems could result!!
Do not proceed with reassignment if the partner is in takeover mode.
Abort reassignment (y/n)? n
```

Enter `n` when asked to abort disk reassignment.

When you are asked to abort disk reassignment, you must answer a series of prompts as shown in the following steps:

- a. The system displays the following message:

```
After the node becomes operational, you must perform a takeover and
giveback of the HA partner node to ensure disk reassignment is
successful.
Do you want to continue (y/n)? y
```

- b. Enter `y` to continue.

The system displays the following message:

```
Disk ownership will be updated on all disks previously belonging to
Filer with sysid <sysid>.
Do you want to continue (y/n)? y
```

- c. Enter `y` to allow disk ownership to be updated.

22. If you are upgrading from a system with external disks to a system that supports internal and external disks (A800 systems, for example), set `node4` as root to confirm that it boots from the root aggregate of `node2`.



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

```
aggr status -r
```

- b. Check the overall status of the `node2` aggregate:

```
aggr status
```

- c. If necessary, bring the `node2` aggregate online:

```
aggr_online root_aggr_from_node2
```

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

```
aggr offline root_aggr_on_node4
```

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

```
aggr options aggr_from_node2 root
```

23. Verify that the controller and chassis are configured as `ha` by entering the following command and observing the output:

```
ha-config show
```

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

```
*> ha-config show
    Chassis HA configuration: ha
    Controller HA configuration: ha
```

Systems record in a PROM whether they are in an HA pair or a 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 correct the configuration:

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

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

24. Destroy the mailboxes on node4:

```
mailbox destroy local
```

25. Exit Maintenance mode:

```
halt
```

The system stops at the boot environment prompt.

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

```
date
```

27. On node4, check the date at the boot environment prompt:

```
show date
```

28. If necessary, set the date on node4:

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

29. On node4, check the time at the boot environment prompt:

```
show time
```

30. If necessary, set the time on node4:

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

31. Verify the partner system ID is set correctly as noted in [Step 19](#) under option.

```
printenv partner-sysid
```

32. If necessary, set the partner system ID on node4:

```
setenv partner-sysid node3_sysid
```

a. Save the settings:

```
saveenv
```

33. Enter the boot menu at the boot environment prompt:

```
boot_ontap menu
```

34. At the boot menu, select option **(6) Update flash from backup config** by entering 6 at the prompt.

The system displays the following message:

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

35. Enter *y* at the prompt.

The boot proceeds normally, and the system prompts you to confirm the system ID mismatch.



The system might reboot twice before displaying the mismatch warning.

36. Confirm the mismatch.

The node might complete one round of rebooting before booting normally.

37. Log in to node4.

Set the FC or UTA/UTA2 configuration on node4

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

About this task

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

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

Configure FC ports on node4

If node4 has FC ports, either onboard or on an FC adapter, you must set port configurations on the node before you bring it into service because the ports are not preconfigured. If the ports are not configured, you might experience a disruption in service.

Before you begin

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

About this task

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



Enter the commands in this section at the cluster 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 the new nodes with the settings that you captured earlier from the original node.

3. Modify the FC ports on node4 as needed:

- To program target ports:

```
system node hardware unified-connect modify -type \ | -t target -adapter  
port_name
```

- To program initiator ports:

```
system node unified-connect modify type \ | -t initiator -adapter port_name
```

-type is the FC4 type, target or initiator.

4. Verify the new settings by entering the following command and examining the output:

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

5. Take one of the following actions:

If the default FC settings on the new nodes are...	Then...
The same as the ones you that captured on the original nodes	Go to Step 9 .
Different from the ones that you captured on the original nodes	Go to Step6 .

6. Exit Maintenance mode:

```
halt
```

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

8. Boot node4 by entering the following command at the boot environment prompt:

```
boot_ontap
```

9. Take one of the following actions:

- Go to [Check and configure UTA/UTA2 ports on node4](#) if node4 has a UTA/UTA2A card or UTA/UTA2 onboard ports.
- Skip the section and go to [Map ports from node2 to node4](#) if node4 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports.

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 enables 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 can check the configuration of the UTA/UTA2 ports on node4 and change it, if necessary.



Enter the commands in this section at the cluster prompt unless directed to enter Maintenance mode. If you have a MetroCluster FC system, you must be in Maintenance mode to configure UTA/UTA2 ports.

Steps

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

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

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

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

3. Examine the output of the `system node hardware unified-connect show` or `ucadmin show` command and determine whether the UTA/UTA2 ports have the personality you want.

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 12 and go to Step 13 .

5. If the system has storage disks and is running Data ONTAP 8.3, boot node4 and enter maintenance mode:

```
boot_ontap maint
```

6. Verify the settings by entering the following command and examining its output:

```
ucadmin show
```

7. Take one of the following actions:

If you are configuring...	Then...
Ports on a UTA/UTA2A card	Go to Step 8 .
Onboard UTA/UTA2 ports	Skip Step 8 and go to Step 9 .

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

```
storage disable adapter adapter_name
```

Adapters in target mode are automatically offline in Maintenance mode.

9. If the current configuration does not match the desired use, enter the following command to 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 the FC target for SAN clients.

10. If the system has storage disks, enter the following command:

```
halt
```

The system stops at the boot environment prompt.

- a. Enter the following command:

```
boot_ontap
```

11. If the system has storage disks, enter the following command:

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

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`.

```
cluster1::> system node hardware unified-connect show
Node      Adapter  Current Mode    Current Type    Pending Mode    Pending Type    Admin Status
-----  -
f-a      1a      fc      initiator -      -      -      online
f-a      1b      fc      target  -      initiator online
f-a      2a      fc      target  cna    -      online
f-a      2b      fc      target  cna    -      online
4 entries were displayed.
```

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

```
network fcp adapter modify -node node_name -adapter adapter_name -state up
```

13. Cable the port.

Map ports from node2 to node4

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

Before you begin

You must already have information about the ports on the new nodes, to access this information refer to [References](#) to link to the *Hardware Universe*. You 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.

About this task

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 may not automatically rejoin 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 this table:

LIF	Node2 ports	Node2 IPspaces	Node2 broadcast domains	Node4 ports	Node4 IPspaces	Node4 broadcast domains
Cluster 1						
Cluster 2						
Cluster 3						
Cluster 4						
Cluster 5						
Cluster 6						
Node management						
Cluster management						
Data 1						
Data 2						
Data 3						
Data 4						
SAN						
Intercluster port						

See the "Recording node2 information" section for the steps to obtain this information.

2. Record all the cabling information for node4, the ports, broadcast domains, and IPspaces, in the previous table using the same procedure in the [Record node2 information](#) section for the steps to obtain this information.
3. Follow these steps to verify if the setup is a two-node switchless cluster:
 - a. Set the privilege level to advanced:

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

```
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. Get node4 into quorum by performing the following steps:

- 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     e1a Cluster
```

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

```
broadcast-domain add-ports -ip-space Cluster -broadcast-domain Cluster -ports
node:port
```

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

```
network port modify -node -port -ip-space Cluster -mtu 9000
```

This example adds Cluster port "e1b" on node4:

```
network port modify -node node4 -port e1b -ip-space Cluster -mtu 9000
```



For a MetroCluster configuration, you might not be able to change the broadcast domain of a port because it is associated with a port hosting the LIF of a sync-destination SVM and see errors similar to, but not restricted to, the following:

```
command failed: This operation is not permitted on a Vserver that is
configured as the destination of a MetroCluster Vserver relationship.
```

Enter the following command from the corresponding sync-source SVM on the remote site to reallocate the sync-destination LIF to an appropriate port:

```
metrocluster vserver resync -vserver vserver_name
```

- 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/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 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_name -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-domain-wide -failover-group fg1
```

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

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.

```

- 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.

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.

Steps

- At the system prompt, log in to node4.
- Verify that node4 is both part of the same cluster as node3 and healthy:

```
cluster show
```

- Verify that node4 can communicate with node3 and that all LIFs are up:

```
network interface show -curr-node node4
```

- Take one of the following actions:

If node4 is...	Then...
In a chassis separate from node3	Connect the HA interconnect between the nodes by completing the following steps: <ol style="list-style-type: none"> Connect the top interconnect port of node3 to the top interconnect port of node4. Connect the bottom interconnect port of node3 to the bottom interconnect port of node4. Go to Step 5.

If node4 is...	Then...
In the same chassis as node3	Go to Step 5 . You do not need to manually connect the HA interconnect between the nodes; in same-chassis configurations, the HA interconnect is connected automatically through the backplane.

5. Take one of the following actions:

If the cluster is...	Then...
In a SAN environment	Complete Step 6 and go to the section Move NAS data LIFs owned by node2 from node3 to node4 and verify SAN LIFs on node4 .
Not in a SAN environment	Skip Step 6 go to the section Move NAS data LIFs owned by node2 from node3 to node4 and verify SAN LIFs on node4 .

6. Verify that both node3 and node4 are in quorum by entering the following command on one of the nodes:

```
event log show -messagename scsiblade.*
```

The following example shows the output when the nodes in the cluster are in quorum:

```
cluster::> event log show -messagename scsiblade.*
Time                Node    Severity    Event
-----
8/13/2012 14:03:51  node1    INFORMATIONAL scsiblade.in.quorum: The scsi-
blade ...
8/13/2012 14:03:51  node2    INFORMATIONAL scsiblade.in.quorum: The scsi-
blade ...
8/13/2012 14:03:48  node3    INFORMATIONAL scsiblade.in.quorum: The scsi-
blade ...
8/13/2012 14:03:43  node4    INFORMATIONAL scsiblade.in.quorum: The scsi-
blade ...
```

Move NAS data LIFs owned by node2 from node3 to node4 and verify SAN LIFs on node4

After you verify the node4 installation and before you relocate node2 aggregates from node3 to node4, you must move the NAS data LIFs owned by node2 currently on node3 from node3 to node4. You also need to verify the SAN LIFs 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 verify that the LIFs are healthy and located on appropriate ports after you bring node4 online.

Steps

1. List all the NAS data LIFs that are not owned by node3 by entering the following command on either node and capturing the output:

```
network interface show -role data -curr-node node3 -is-home false
```

2. If the cluster is configured for SAN LIFs, record the SAN LIFs and existing configuration information in this [worksheet](#) for use later in the procedure.
 - a. List the SAN LIFs on node3 and examine the output:

```
network interface show -data-protocol fc*
```

The system returns output similar to the following example:

```
cluster1::> net int show -data-protocol fc*
(network interface show)
Current Is      Logical      Status      Network      Current
Vserver        Interface    Admin/Oper  Address/Mask  Node
Port           Home
-----
-----
svm2_cluster1
      lif_svm2_cluster1_340
                        up/up      20:02:00:50:56:b0:39:99
                                                cluster1-01
1b      true
      lif_svm2_cluster1_398
                        up/up      20:03:00:50:56:b0:39:99
                                                cluster1-02
1a      true
      lif_svm2_cluster1_691
                        up/up      20:01:00:50:56:b0:39:99
                                                cluster1-01
1a      true
      lif_svm2_cluster1_925
                        up/up      20:04:00:50:56:b0:39:99
                                                cluster1-02
1b      true
4 entries were displayed.
```

- b. List the existing configurations and examine the output:

```
fcv adapter show -fields switch-port,fc-wwpn
```

The system returns output similar to the following example:

```

cluster1::> fcp adapter show -fields switch-port,fc-wwpn
(network fcp adapter show)
node          adapter  fc-wwpn                               switch-port
-----
cluster1-01  0a       50:0a:09:82:9c:13:38:00             ACME Switch:0
cluster1-01  0b       50:0a:09:82:9c:13:38:01             ACME Switch:1
cluster1-01  0c       50:0a:09:82:9c:13:38:02             ACME Switch:2
cluster1-01  0d       50:0a:09:82:9c:13:38:03             ACME Switch:3
cluster1-01  0e       50:0a:09:82:9c:13:38:04             ACME Switch:4
cluster1-01  0f       50:0a:09:82:9c:13:38:05             ACME Switch:5
cluster1-01  1a       50:0a:09:82:9c:13:38:06             ACME Switch:6
cluster1-01  1b       50:0a:09:82:9c:13:38:07             ACME Switch:7
cluster1-02  0a       50:0a:09:82:9c:6c:36:00             ACME Switch:0
cluster1-02  0b       50:0a:09:82:9c:6c:36:01             ACME Switch:1
cluster1-02  0c       50:0a:09:82:9c:6c:36:02             ACME Switch:2
cluster1-02  0d       50:0a:09:82:9c:6c:36:03             ACME Switch:3
cluster1-02  0e       50:0a:09:82:9c:6c:36:04             ACME Switch:4
cluster1-02  0f       50:0a:09:82:9c:6c:36:05             ACME Switch:5
cluster1-02  1a       50:0a:09:82:9c:6c:36:06             ACME Switch:6
cluster1-02  1b       50:0a:09:82:9c:6c:36:07             ACME Switch:7
16 entries were displayed

```

3. Take one of the following actions:

If node2...	Description
Had interface groups or VLANs configured	Go to Step 4 .
Did not have interface groups or VLANs configured	Skip Step 4 and go to Step 5 .

4. Take the following steps to migrate any NAS data LIFs hosted on interface groups and VLANs that originally were on node2 from node3 to node4.

- a. Migrate any LIFs hosted on node3 that previously belonging to node2 on an interface group to a port on node4 that is capable of hosting LIFs on the same network by entering the following command, once for each LIF:

```

network interface migrate -vserver vservice_name -lif lif_name -destination
-node node4 -destination-port netport|ifgrp

```

- b. Modify the home port and home node of the LIFs in [Substep a](#) to the port and node currently hosting the LIFs by entering the following command, once for each LIF:

```

network interface modify -vserver vservice_name -lif datalif_name -home-node
node4 home-port netport|ifgrp

```

- c. Migrate any LIFs hosted on node3 that previously belonged to node2 on a VLAN port to a port on node4 that is capable of hosting LIFs on the same network by entering the following command, once

for each LIF:

```
network interface migrate -vserver vserver_name -lif datalif_name
-destination-node node4 -destination-port netport|ifgrp
```

- d. Modify the home port and home node of the LIFs in [Substep c](#) to the port and node currently hosting the LIFs by entering the following command, once for each LIF:

```
network interface modify -vserver vserver_name -lif datalif_name -home-node
node4 home-port netport|ifgrp
```

5. Take one of the following actions:

If the cluster is configured for...	Then...
NAS	Complete Step 6 through Step 9 , skip Step 10 , and complete Step 11 through Step 14 .
SAN	Skip Step 6 through Step 9 , and complete Step 10 through Step 14 .
Both NAS and SAN	Complete Step 6 through Step 14 .

6. If you have data ports that are not the same on your platforms, enter the following command to add the ports to the broadcast domain:

```
network port broadcast-domain add-ports -ipspace IPspace_name -broadcast
-domain mgmt ports node:port
```

The following example adds port "e0a" on node "6280-1" and port "e0i" on node "8060-1" to broadcast domain mgmt in the IPspace Default:

```
cluster::> network port broadcast-domain add-ports -ipspace Default
-broadcast-domain mgmt -ports 6280-1:e0a, 8060-1:e0i
```

7. Migrate each NAS data LIF to node4 by entering the following command, once for each LIF:

```
network interface migrate -vserver vserver_name -lif datalif_name -destination
-node node4 -destination-port netport|ifgrp -home-node node4
```

8. Make sure that the data migration is persistent:

```
network interface modify -vserver vserver_name -lif datalif_name -home-port
netport|ifgrp
```

9. Verify the status of all links as up by entering the following command to list all the network ports and examining its output:

```
network port show
```

The following example shows the output of the `network port show` command with some LIFs up and others down:

```

cluster::> network port show

```

						Speed	
(Mbps)							
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper	

node3							
	a0a	Default	-	up	1500	auto/1000	
	e0M	Default	172.17.178.19/24	up	1500	auto/100	
	e0a	Default	-	up	1500	auto/1000	
	e0a-1	Default	172.17.178.19/24	up	1500	auto/1000	
	e0b	Default	-	up	1500	auto/1000	
	e1a	Cluster	Cluster	up	9000	auto/10000	
	e1b	Cluster	Cluster	up	9000	auto/10000	
node4							
	e0M	Default	172.17.178.19/24	up	1500	auto/100	
	e0a	Default	172.17.178.19/24	up	1500	auto/1000	
	e0b	Default	-	up	1500	auto/1000	
	e1a	Cluster	Cluster	up	9000	auto/10000	
	e1b	Cluster	Cluster	up	9000	auto/10000	
12 entries were displayed.							

10. If the output of the `network port show` command displays network ports that are not available in the new node and are present in the old nodes, delete the old network ports by completing the following substeps:

- a. Enter the advanced privilege level by entering the following command:

```
set -privilege advanced
```

- b. Enter the following command, once for each old network port:

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

- c. Return to the admin level by entering the following command:

```
set -privilege admin
```

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

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

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

The system returns output similar to the following example:

```

cluster::> network interface show -data-protocol iscsi|fc -home-node
node4

```

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			
vs0	a0a	up/down	10.63.0.53/24	node4
a0a	true			
	data1	up/up	10.63.0.50/18	node4
e0c	true			
	rads1	up/up	10.63.0.51/18	node4
e1a	true			
	rads2	up/down	10.63.0.52/24	node4
e1b	true			
vs1				
	lif1	up/up	172.17.176.120/24	node4
e0c	true			
	lif2	up/up	172.17.176.121/24	node4

- b. Verify that the new adapter and switch-port configurations are correct by comparing the output from the `fc adapter show` command with the new configuration information that you recorded in the worksheet in [Step 2](#).

List the new SAN LIF configurations on node4:

```
fc adapter show -fields switch-port,fc-wwpn
```

The system returns output similar to the following example:

```

cluster1::> fcp adapter show -fields switch-port,fc-wwpn
(network fcp adapter show)
node          adapter  fc-wwpn          switch-port
-----
cluster1-01  0a       50:0a:09:82:9c:13:38:00  ACME Switch:0
cluster1-01  0b       50:0a:09:82:9c:13:38:01  ACME Switch:1
cluster1-01  0c       50:0a:09:82:9c:13:38:02  ACME Switch:2
cluster1-01  0d       50:0a:09:82:9c:13:38:03  ACME Switch:3
cluster1-01  0e       50:0a:09:82:9c:13:38:04  ACME Switch:4
cluster1-01  0f       50:0a:09:82:9c:13:38:05  ACME Switch:5
cluster1-01  1a       50:0a:09:82:9c:13:38:06  ACME Switch:6
cluster1-01  1b       50:0a:09:82:9c:13:38:07  ACME Switch:7
cluster1-02  0a       50:0a:09:82:9c:6c:36:00  ACME Switch:0
cluster1-02  0b       50:0a:09:82:9c:6c:36:01  ACME Switch:1
cluster1-02  0c       50:0a:09:82:9c:6c:36:02  ACME Switch:2
cluster1-02  0d       50:0a:09:82:9c:6c:36:03  ACME Switch:3
cluster1-02  0e       50:0a:09:82:9c:6c:36:04  ACME Switch:4
cluster1-02  0f       50:0a:09:82:9c:6c:36:05  ACME Switch:5
cluster1-02  1a       50:0a:09:82:9c:6c:36:06  ACME Switch:6
cluster1-02  1b       50:0a:09:82:9c:6c:36:07  ACME Switch:7
16 entries were displayed

```



If a SAN LIF in the new configuration is not on an adapter that is still attached to the same switch-port, it might cause a system outage when you reboot the node.

c. If node4 has any SAN LIFs or groups of SAN LIFs that are on a port that did not exist on node2, move them to an appropriate port on node4 by entering one of the following commands:

i. Set the LIF status to down:

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

ii. Remove the LIF from the port set:

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

iii. Enter one of the following commands:

- Move a single LIF:

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

- Move all the LIFs on a single nonexistent or incorrect port to a new port:

```
network interface modify {-home-port port_on_node2 -home-node node2
-role data} -home-port new_home_port_on_node4
```


Relocate node2 non-root aggregates from node3 to node4

Having relocated node2's non-root aggregates to node3, you now must relocate them from node3 to node4.

Steps

1. Enter the following command on either controller, and examine the output to identify which non-root aggregates to relocate:

```
storage aggregate show -owner-name node3 -home-id node2_system_id
```

2. Relocate the aggregates by completing the following substeps:

- a. Access the advanced privilege level by entering the following command on either node:

```
set -privilege advanced
```

- b. Enter the following command:

```
storage aggregate relocation start -node node3 -destination node4 -aggregate  
-list aggr_name1, aggr_name2... -ndo-controller-upgrade true
```

The aggregate list is the list of aggregates owned by node4 that you obtained in [Step 1](#).

- c. When prompted, enter *y*.

Relocation occurs in the background. It could take anywhere from a few seconds to a couple of minutes to relocate an aggregate. The time includes both client outage and non-outage portions. The command does not relocate any offline or restricted aggregates.

- d. Return to the admin level:

```
set -privilege admin
```

3. Check the relocation status:

```
storage aggregate relocation show -node node3
```

The output will display *Done* for an aggregate after it has been relocated.



Wait until all the node2 aggregates have been relocated to node4 before proceeding to the next step.

4. Take one of the following actions:

If relocation of...	Then...
All aggregates was successful	Go to Step 5 .

If relocation of...	Then...
Any aggregates failed, or were vetoed	<p>a. Check the EMS logs for the corrective action.</p> <p>b. Perform the corrective action.</p> <p>c. Access the advanced privilege level by entering the following command on either node:</p> <pre>set -privilege advanced</pre> <p>d. Relocate any failed or vetoed aggregates:</p> <pre>storage aggregate relocation start -node node3 destination node4 -aggregate-list aggr_name1, aggr_name2... ndo-controller-upgrade true</pre> <p>The aggregate list is the list of failed or vetoed aggregates.</p> <p>e. When prompted, enter <i>y</i>.</p> <p>f. Return to the admin level by entering the following command:</p> <pre>set -privilege admin</pre> <p>If necessary, you can force the relocation using one of the following methods:</p> <ul style="list-style-type: none"> • Overriding veto checks: <pre>storage aggregate relocation start -override -vetoes -ndo-controller-upgrade</pre> <ul style="list-style-type: none"> • Overriding destination checks: <pre>storage aggregate relocation start -override -destination-checks -ndocontroller-upgrade</pre> <p>For more information about storage aggregate relocation commands refer to References to link to <i>Disk and aggregate management with the CLI</i> and the <i>ONTAP 9 Commands: Manual Page Reference</i>.</p>

5. Verify that all node2 non-root aggregates are online and their state on node4:

```
storage aggregate show -node node4 -state offline -root false
```

The node2 aggregates were listed in the output of the command in [Step 1](#).

6. If any aggregate has gone offline or become foreign, bring it online by using the following command for each aggregate:

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

7. Verify that all the volumes in node2 aggregates are online on node4:

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

8. If any volumes are offline on node4, bring them online:

```
volume online -vserver vserver-name -volume volume_name
```

9. Send a post-upgrade AutoSupport message to NetApp for node4:

```
system node autosupport invoke -node node4 -type all -message "node2  
successfully upgraded from platform_old to platform_new"
```

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.