



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
Verify the node4 installation	11
Restore network configuration on node4	12
Restore key-manager configuration on node4	17
Move non-root aggregates and NAS data LIFs owned by node2 from node3 to node4	18

Stage 5. Install and boot node4

Install and boot node4

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

About this task

The relocation operation is paused at the beginning of this stage. This process is mostly automated; the operation pauses to enable you to check its status. You must manually resume the operation.

You 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](#).

Steps

1. Make sure that node4 has sufficient rack space.

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

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

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

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

- Console (remote management port)
- Cluster and HA ports
- Data ports
- Cluster and node management ports
- Serial-attached SCSI (SAS) and Ethernet storage ports
- SAN configurations: iSCSI Ethernet, FC, and NVMe/FC switch ports

You might need to change the interconnect cables between the old and new controllers to allow interoperability between the different controller and card models. Refer to the [system installation procedures](#) for a cabling map of the Ethernet storage shelves for your systems.



For controllers introduced in ONTAP 9.15.1 and later, cluster and HA interconnects use the same ports. For switch connected configurations, it is required to connect similar ports to the same cluster switches. For example, when upgrading to an AFF A1K from an existing controller, you should connect e1a ports on both nodes to one switch and e7a ports on both nodes to the second switch.

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:

- 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 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 using the following command at the boot environment prompt: <code>ifconfig e0M -auto</code>

If Dynamic Host Configuration Protocol (DHCP) is...	Then...
Not running	<p>Manually configure the connection by entering the following command at the boot environment prompt:</p> <pre>ifconfig e0M -addr=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 DNS domain name (optional).</p> <p> 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:

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

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



Do not interrupt the boot.

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

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

Disregard the following message:

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

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



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

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

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

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

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

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

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

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

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

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

```
ha-config show
```

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

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



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

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

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

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

14. Confirm that all Ethernet ports used to connect to the Ethernet shelves are configured as storage:

```
storage port show
```

The output displayed depends on the system configuration. The following output example is for a node with a single storage card in slot11. The output for your system might be different:

```
*> storage port show  
Port Type Mode Speed(Gb/s) State Status VLAN ID  
---- ---- ---- ---- ---- ---- ----  
e11a ENET storage 100 Gb/s enabled online 30  
e11b ENET storage 100 Gb/s enabled online 30
```

15. Modify the ports that are not set to storage:

```
storage port modify -p <port> -m storage
```

All Ethernet ports connected to storage shelves must be configured as storage to allow access to the disks and shelves.

16. Exit maintenance mode:

```
halt
```

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

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

```
date
```

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

```
show date
```

19. If necessary, set the date on node4:

```
set date <mm/dd/yyyy>
```

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

```
show time
```

21. If necessary, set the time on node4:

```
set time <hh:mm:ss>
```

22. 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
```

23. Verify the partner-sysid for node4:

```
printenv partner-sysid
```

24. 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>

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

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

25. Boot the node into boot menu:

`boot_ontap menu`

26. When node4 stops at the boot menu, reassign the node2 disks to node4 by running the following command on 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
.
.
.
<output truncated>
.
All rights reserved.
*****
*                               *
* Press Ctrl-C for Boot Menu. *
*                               *
*****
.

<output truncated>
.

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

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

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

Normal Boot is prohibited.

Please choose one of the following:

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

Selection (1-11)? boot_after_controller_replacement

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

.

.

<output truncated>

.

.

Controller Replacement: Provide name of the node you would like to replace:

<nodename of the node being replaced>

Controller Replacement: Provide High Availability partner of node1:

<nodename of the partner of the node being replaced>

Changing sysid of node node2 disks.

Fetched sanown old_owner_sysid = 536940063 and calculated old sys id = 536940063

Partner sysid = 4294967295, owner sysid = 536940063

.

.

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



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

27. 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:

If the root volume is encrypted, recover the onboard key-management secrets so the system can find 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 27](#) until the system boots normally.

Verify the node4 installation

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

About this task

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

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

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

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

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

Steps

1. Verify that node4 has joined quorum:

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

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

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

```
cluster show
```

3. Switch to advanced privilege mode:

```
set advanced
```

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

```
system controller replace show
```

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

5. Resume the controller replacement operation:

```
system controller replace resume
```

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

```

Cluster::*> system controller replace show
Node          Status          Error-Action
-----
-----
Node2 (now node4) Paused-for-intervention      Follow the instructions
given in
                                                Step Details
Node2

Step Details:
-----
To complete the Network Reachability task, the ONTAP network
configuration must be
manually adjusted to match the new physical network configuration of the
hardware.
This includes:
1. Re-create the interface group, if needed, before restoring VLANs. For
detailed
commands and instructions, refer to the "Re-creating VLANs, ifgrps, and
broadcast
domains" section of the upgrade controller hardware guide for the ONTAP
version
running on the new controllers.
2. Run the command "cluster controller-replacement network displaced-
vlans show"
to check if any VLAN is displaced.
3. If any VLAN is displaced, run the command "cluster controller-
replacement
network displaced-vlans restore" to restore the VLAN on the desired
port.
2 entries were displayed.

```



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

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

Restore network configuration on node4

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

About this task

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

Steps

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

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

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

2. List the broadcast domains on the cluster:

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

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

```
network port reachability show
```

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

```

ClusterA::*> network port reachability show
Node      Port      Expected Reachability          Reachability
Status

-----
-----
node1_node3
    e0M      Default:Mgmt          ok
    e10a     Default:Default       ok
    e10b      -                  no-reachability
    e10c     Default:Default       ok
    e10d      -                  no-reachability
    e1a      Cluster:Cluster      ok
    e1b      -                  no-reachability
    e7a      Cluster:Cluster      ok
    e7b      -                  no-reachability

node2_node4
    e0M      Default:Mgmt          ok
    e10a     Default:Default       ok
    e10b      -                  no-reachability
    e10c     Default:Default       ok
    e10d      -                  no-reachability
    e1a      Cluster:Cluster      ok
    e1b      -                  no-reachability
    e7a      Cluster:Cluster      ok
    e7b      -                  no-reachability

18 entries were displayed.

```

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

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

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

The output looks like the following example:

```
Cluster ::> reachability repair -node node2_node4 -port e10a
```

Warning: Repairing port "node2_node4: e10a" may cause it to move into a different broadcast domain, which can cause LIFs to be re-homed away from the port. Are you sure you want to continue? {y|n}:

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

from the reachability status of the broadcast domain where it is currently located.

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

Verify that all physical ports have their expected reachability:

```
network port reachability show
```

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

5. If interface group configuration does not match the new controller physical port layout, modify it by using the following steps.

a. You must first remove physical ports that should be interface group member ports from their broadcast domain membership. You can do this by using the following command:

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

b. Add a member port to an interface group:

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

c. The interface group is automatically added to the broadcast domain about a minute after the first member port is added.

d. Verify that the interface group was added to the appropriate broadcast domain:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

a. List displaced VLANs:

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

Output like the following should display:

```
Cluster::*:> displaced-vlans show
(cluster controller-replacement network displaced-vlans show)
      Original
  Node      Base Port      VLANs
  -----  -----
Node1        a0a        822, 823
            e10a        822, 823
```

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

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

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

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

The following is an example of restoring displaced VLANs on port "e10a" to "e10b":

```
Cluster::*:> displaced-vlans restore -node node2_node4 -port e10a
-destination-port e10b
```

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

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

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

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

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

10. Verify port reachability:

```
network port reachability show
```

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

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

```
network port show
```

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

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

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

a. List any LIFs that are displaced:

```
displaced-interface show
```

b. Restore LIF home ports:

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

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

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

Restore key-manager configuration on node4

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

About this task

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

Steps

1. Run the following command from node4:

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

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

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

Example

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

node      vserver      key-server      key-id
restored

-----
-----
node4      svm1          ""          0000000000000000200000000000a008a81976
true                                2190178f9350e071fbb90f000000000000000000
```

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

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

About this task

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

Steps

1. The iSCSI LIFs automatically find the correct home ports using the reachability scan. The FC and NVMe/FC SAN LIFs do not move automatically. They continue to show the home port they were on before upgrading.

Check the SAN LIFs on node4:

- a. Modify any iSCSI SAN LIFs reporting a "down" operation status to the new data ports:

```
network interface modify -vserver <vserver> -lif <iscsi_san_lif>  
admin down
```

```
network interface modify -vserver <vserver> -lif <iscsi_san_lif> port  
<new_port> node <node>
```

```
network interface modify -vserver <vserver> -lif <iscsi_san_lif>
```

b. Modify any FC and NVMe/FC SAN LIFs that are home to the new controller and reporting a "down" operational status to the FCP ports on the new controller:

```
network interface modify -vserver <vserver> -lif <fc_san_lif> admin  
down
```

```
network interface modify -vserver <vserver> -lif <fc_san_lif> port  
<new_port> node <node>
```

```
network interface modify -vserver <vserver> -lif <fc_san_lif>
```

2. 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.

3. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following checks:

- Cluster health check
- Cluster LIF status check

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

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

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

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

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

5. If necessary, restore and revert displaced LIFs or manually migrate and modify the node2 LIFs that failed to relocate automatically to node4.

Restore and revert displaced LIFs

- List any displaced LIFs:

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

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

```
cluster controller-replacement network displaced-interface
restore-home-node -node <node4_nodename> -vserver <vserver name>
-lif-name <lif_name>
```

Manually migrate and modify LIFs

- Migrate the LIFs that failed to relocate automatically to node4:

```
network interface migrate -vserver <vserver name> -lif <lif_name>
-destination-node <node4_nodename> -destination-port
<port_on_node4>
```

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

```
network interface modify -vserver <vserver_name> -lif
<data_lif_name> -home-node <node4_nodename> -home-port
<home_port>
```

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

```
system controller replace resume
```

The system performs the following post-checks:

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

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.