Stage 6. Boot node2 with the replacement system modules

Upgrade controllers

NetApp
February 10, 2026

Table of Contents

Stage 6. Boot node2 with the replacement system modules

Cable node2 for shared cluster-HA and storage

If you are performing one of the following upgrades, you need to connect the cluster, HA, storage, data, and management connections that were previously connected to the node2 on the existing system to the newly installed node2 on the replacement system.

Existing system	Replacement system
AFF A250	AFF A30, AFF A50
AFF C250	AFF C30, AFF C60
AFF A800	AFF A70, AFF A90
AFF C800	AFF C80

Connect the e0M and BMC ports

If the existing system has a management port (e0M) and a BMC port, the e0M and BMC ports are combined and accessed through the "wrench" port on the replacement system. You must ensure that the e0M and BMC ports are connected to the same switch and subnet on the existing system before connecting to the replacement system.

If the…	Then…
e0M and BMC IP addresses are on the same IP subnet	Connect either the e0M or BMC port on the existing system to the "wrench" port on the replacement system.
e0M and BMC IP addresses are on different subnets	1. Merge the e0M and BMC IP addresses into one IP subnet. 2. Connect either the e0M or BMC port on the existing system to the "wrench" port on the replacement system.

Connect to a two-node switchless cluster

The following tables show the switch port usage for two-node switchless cluster configurations.

Port type	AFF A800, AFF C800	AFF A90	AFF A70, AFF C80
Cluster	e0a	e1a	e1a
Cluster	e1a	e7a (Use e1b if there is no e7a)	e1b
HA	e0b	Don’t connect	Don’t connect
HA	e1b	Don’t connect	Don’t connect

Port type	AFF A800, AFF C800	AFF A90	AFF A70, AFF C80
SAS storage ports (if present and used)	Any available port	Any available port	Any available port
Ethernet storage ports for NS224 shelves	Any available port	Refer to Ethernet storage connectivity mapping	Refer to Ethernet storage connectivity mapping

Port type	AFF A250, AFF C250	AFF A30, AFF C60	AFF A50
Cluster	e0c	e1a (Use e1a for temporary cluster interconnect)	e1a (Use e1a for temporary cluster interconnect)
Cluster	e0d	e1b (Use e1b for temporary cluster interconnect)	e1b (Use e1b for temporary cluster interconnect)
HA	e0c HA port is shared with Cluster port	e4a on node1 is directly connected to e4a on node2 using a 100 GbE cable	e4a on node1 is directly connected to e4a on node2 using a 100 GbE cable
HA	e0d HA port is shared with Cluster port	e2a on node1 is directly connected to e2a on node2 using a 100 GbE cable If e2a isn’t present or doesn’t support 100 GbE, directly connect e4b on node1 to e4b on node2 using a 100 GbE cable.	e2a on node1 directly connected to e2a on node2 using a 100 GbE cable If e2a isn’t present or doesn’t support 100 GbE, directly connect e4b on node1 to e4b on node2 using a 100 GbE cable.
Ethernet storage port	Any available port	e3a, e3b	e3a, e3b
SAS storage port	Any available port	3a, 3b	3a, 3b

Connect to a switch-attached cluster

For a switch-attached cluster, check that you meet the following requirements for the AFF A30, AFF A50, AFF A70, AFF A90, AFF C30, AFF C60, or AFF C80 (replacement) node:

- The identical cluster ports on the replacement node are on the same switch. For example, on completion of the upgrade, e1a on node1 and e1a on node2 should be attached to one cluster switch. Similarly, the second cluster port from both nodes should be attached to the second cluster switch. Cross-connection of shared cluster-HA ports, where e1a from node1 is connected to switchA and e1a from node2 is connected to switchB, results in HA communication failures.
- The replacement node uses shared cluster-HA Ethernet ports.
- Verify that the cluster switches are installed with a reference configuration file (RCF) that supports shared cluster-HA ports:
    1. Remove the existing configuration on the switch:

If your switch model is…	Go to…
Cisco Nexus	The Knowledge Base article How to clear configuration on a Cisco interconnect switch while retaining remote connectivity
Broadcom BES-53248	The Knowledge Base article How to clear configuration on a Broadcom interconnect switch while retaining remote connectivity

2. Configure and verify the switch setup:

If your switch model is…	Go to…
Cisco Nexus 9336C-FX2	Upgrade your Reference Configuration File (RCF)
Broadcom BES-53248	Upgrade the Reference Configuration File (RCF)
NVIDIA SN2100	Install or upgrade the Reference Configuration File (RCF) script

If the cluster switch only supports 10/25 GbE speeds, you must use an X60130A, 4-port 10/25GbE card in slot1 or slot2 on the replacement system for cluster interconnect.

Boot node2 with the replacement system modules

Node2 with the replacement modules is now ready to boot. The supported replacement modules are listed in the supported systems matrix.

You only move the console and management connections when you upgrade by swapping the system modules.

Steps

1. 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	`setenv bootarg.storageencryption.support` **`true`**
NetApp non-FIPS SEDs	`setenv bootarg.storageencryption.support` **`false`**

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

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

`secrets`.

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

2. Boot the node into the boot menu:

   `boot_ontap menu`

3. When the node stops at the boot menu, reassign the old node2 disks to the replacement node2 by running the following command on node2:

   `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.

   > If `[localhost:disk.encryptNoSupport:ALERT]: Detected FIPS-certified encrypting drive` and, or, `[localhost:diskown.errorDuringIO:error]: error 3 (disk failed) on disk` errors occur, perform the following steps:
   >
   > 1. Halt the node at the LOADER prompt.
   > 2. Check and reset the storage encryption bootargs mentioned in Step 1.
   > 3. At the LOADER prompt, boot up:
   >
   >   `boot_ontap`

   You can use the following example as a reference:

Expand the console output example

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

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

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

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



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

Please choose one of the following:

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

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

.
.
<output truncated>
.
.
Controller Replacement: Provide name of the node you would like to
replace:<nodename of the node being replaced>
Controller Replacement: Provide High Availability partner of node1:
<nodename of the partner of the node being replaced>
```

```
Changing sysid of node node1 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:
```

> The system IDs shown in the preceding example are example IDs. The actual system IDs of the nodes that you are upgrading will be different.
>
> Between entering node names at the prompt and the login prompt, the node reboots a few times to restore the environment variables, update firmware on the cards in the system, and for other ONTAP updates.

Verify the node2 installation

You must verify the node2 installation with the replacement system modules. Because there is no change to physical ports, you are not required to map the physical ports from the old node2 to the replacement node2.

About this task

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

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

Steps

1. Verify that node2 has joined quorum:

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

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

2. This step applies to the following upgrade configurations. For all other system upgrades, skip this step and go to Step 3:

   - Two node switchless clusters
   - Switch attached AFF A250 or AFF C250 systems upgrading to an AFF A50, AFF A30, AFF C30, or AFF C60 system.

   If node2 doesn’t join quorum automatically:

   a. Check the IPspace of ports e1a and e1b:

      ```
      network port show
      ```

   b. If the IPspace isn’t "Cluster", change the IPspace to "Cluster" on e1a and e1b:

      ```
      network port modify -node <node_name> -port <port> -ipspace Cluster
      ```

   c. Verify that the IPspace of ports e1a and e1b is "Cluster":

      ```
      network port show
      ```

d. Migrate node2 cluster LIFs to e1a and e1b:

```
network interface migrate -vserver Cluster -lif <cluster_lif1> -destination
-node <node2_name> -destination-port <port_name
```

3. Verify that node2 and node1 are part of the same cluster and that the cluster is healthy:

```
cluster show
```

4. Switch to advanced privilege mode:

```
set advanced
```

5. This step only applies to two-node switchless configuration upgrades from an AFF A250 or AFF C250 to an AFF A50, AFF A30, AFF C60, or AFF C30. For all other system upgrades, skip this step and go to Step 6:

   Verify that e4a, e2a, e1a, e1b ports or e4a, e4b, e1a, e1b ports are the cluster ports in "Cluster" broadcast domain.

   The AFF A50, AFF A30, AFF C30, and AFF C60 systems share cluster and HA ports. You can safely migrate all cluster LIFs to e4a, e4b or e4a, e2a on node1 and node2:

   a. List the home ports and current ports for all cluster LIFs:

   ```
   network interface show -role Cluster -fields home-port,curr-port
   ```

   b. On node1 and node2, migrate the cluster LIFs that are using e1a as the home port to e4a:

   ```
   network interface migrate -vserver Cluster -lif <cluster_lif1> -destination
   -node <node> -destination-port e4a
   ```

   c. On node1 and node2, modify the cluster LIFs migrated in substep b to use e4a as the home port:

   ```
   network interface modify -vserver Cluster -lif <cluster_lif> -home-port e4a
   ```

   d. Verify that the cluster is in quorum:

   ```
   cluster show
   ```

   e. Repeat substep b and substep c to migrate and modify the second cluster LIF on each node to e2a or e4b:

   If e2a is present and is a 100GbE network port, this is the default second cluster port. If e2a isn't a 100GbE network port, ONTAP uses e4b as the second cluster and HA port.

   f. Remove e1a and e1b from "Cluster" broadcast domain:

   ```
   broadcast-domain remove-ports -broadcast-domain Cluster -ipspace Cluster
   -ports <node_name>:e1a
   ```

   g. Verify that only cluster ports e4a, e2a or e4a, e4b are in "Cluster" broadcast domain

   ```
   broadcast domain show
   ```

h. Remove the cable connections between e1a node1 and e1a node2, and e1b node1 and e1b node2 to ensure only valid cluster-HA connections are used and there is no redundant connectivity.

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

```
system controller replace show

system controller replace show-details
```

7. Resume the controller replacement operation:

```
system controller replace resume
```

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

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

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


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

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

9. With the controller replacement in a paused state, proceed to Restore network configuration on node2.

Restore network configuration on node2

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

About this task

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

Steps

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

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

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

2. List the broadcast domains on the cluster:

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

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

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

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

```
Cluster::> reachability show -node node1
  (network port reachability show)
Node      Port     Expected Reachability                Reachability
Status
--------- -------- ------------------------------------
---------------------
Node1
          a0a      Default:Default                      ok
          a0a-822  Default:822                          ok
          a0a-823  Default:823                          ok
          e0M      Default:Mgmt                         ok
          e1a      Cluster:Cluster                      ok
          e1b      -                                    no-reachability
          e2a      -                                    no-reachability
          e2b      -                                    no-reachability
          e3a      -                                    no-reachability
          e3b      -                                    no-reachability
          e7a      Cluster:Cluster                      ok
          e7b      -                                    no-reachability
          e9a      Default:Default                      ok
          e9a-822  Default:822                          ok
          e9a-823  Default:823                          ok
          e9b      Default:Default                      ok
          e9b-822  Default:822                          ok
          e9b-823  Default:823                          ok
          e9c      Default:Default                      ok
          e9d      Default:Default                      ok
20 entries were displayed.
```

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

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

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

   a. Physical ports
   b. VLAN ports

   You should see output like the following example:

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

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

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

Verify that all physical ports have their expected reachability:

```
network port reachability show
```

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

5. Verify port reachability:

   ```
   network port reachability show
   ```

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

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

   ```
   network port show
   ```

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

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

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

   a. List any LIFs that are displaced:

      ```
      displaced-interface show
      ```

   b. Restore LIF home nodes and home ports:

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

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

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

Restore key-manager configuration on node2

If you are using NetApp Aggregate Encryption (NAE) or NetApp Volume Encryption (NVE) to encrypt volumes on the system that you are upgrading, the encryption configuration must be synchronized to the new nodes. If you do not resynchronize the key-manager, when you relocate the node2 aggregates from the upgraded node1 to the upgraded node2 by using ARL, failures might occur because node2 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 node2:

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

2. Verify that the SVM-KEK key is restored to "true" on node2 before you relocate the data aggregates:

   ```
   ::> security key-manager key query -node node2 -fields restored -key
   -type SVM-KEK
   ```

   Example

   ```
   ::> security key-manager key query -node node2 -fields restored -key
   -type SVM-KEK

   node     vserver   key-server   key-id
   restored
   -------- --------- -----------  ----------------------------------------
   --------
   node2    svm1      ""           000000000000000002000000000000a008a81976
   true
                                   2190178f9350e071fbb90f00000000000000000
   ```

Verify the RCF configuration on cluster switches

At this stage in the upgrade procedure, all data aggregates should be on node1. If you're upgrading a configuration with switch-attached clusters, you need to validate that the cluster switch reference configuration file (RCF) supports the shared cluster/HA ports for the new controllers.

If you're upgrading to a two-node switchless cluster configuration, you can skip this section and go to Move non-root aggregates and NAS data LIFs back to node2.

Steps

1. Switch to advanced privilege mode:

   ```
   set advanced
   ```

2. Check the status of "IC RDMA":

   ```
   ha interconnect status show
   ```

   In the output, the "IC RDMA Connection" should have the status `Up`.

If the "IC RDMA Connection" status is …	Then…
`Up`	Go to Move non-root aggregates and NAS data LIFs back to node2.
`Down`	Go to Step 3.

3. Check the cluster ports and switch RCF.

   For more information, see Connect to a switch-attached cluster.

4. Verify that the "IC RDMA Connection" status has changed to `Up`:

   ```
   ha interconnect status show
   ```

What's next

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

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

After verifying the network configuration on node2, you need to relocate the NAS data LIFs owned by node2 from node1 to node2 and confirm that the SAN LIFs exist on node2.

About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. Moving SAN LIFs is not necessary for cluster or service health during the upgrade. SAN LIFs are not moved unless they need to be mapped to new ports.

You verify that the LIFs are healthy and located on the correct ports after you bring node2 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 system pauses the operation at this stage in the network reachability check

2. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following checks:

- Cluster health check
- Cluster LIF status check

After performing these checks, the system relocates the non-root aggregates and NAS data LIFs back to node2, which is now running on the replacement controller.

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

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

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

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

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

Restore and revert displaced LIFs

a. List any displaced LIFs:

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

b. If any LIFs are displaced, restore the home node back to node2:

```
cluster controller-replacement network displaced-interface
restore-home-node -node <node2_nodename> -vserver <vserver name>
-lif-name <lif_name>
```

Manually migrate and modify LIFs

a. Migrate the LIFs that failed to relocate automatically to node2:

```
network interface migrate -vserver <vserver name> -lif <lif_name>
-destination-node <node2_nodename> -destination-port
<port_on_node2>
```

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

```
network interface modify -vserver <vserver_name> -lif
<data_lif_name> -home-node <node2_nodename> -home-port
<home_port>
```

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

```
system controller replace resume
```

The system performs the following post-checks:

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

Copyright information

Trademark information