



Stage 3. Boot node1 with the replacement system modules

Upgrade controllers

NetApp
January 17, 2025

Table of Contents

- Stage 3. Boot node1 with the replacement system modules 1
 - Stage 3 overview 1
 - Cable node1 for shared cluster-HA and storage 1
 - Boot node1 with the replacement system modules 3
 - Verify the node1 installation 8
 - Restore key-manager configuration on the upgraded node1 13
 - Move node1 non-root aggregates and NAS data LIFs from node2 to the upgraded node1 13

Stage 3. Boot node1 with the replacement system modules

Stage 3 overview

During Stage 3, you connect the shared cluster-HA and storage connections for the external shelves, if any, boot node1 with the upgraded system modules, and verify the upgraded node1 installation. If you are using NetApp Volume Encryption (NVE), you restore key-manager configuration. You also relocate node1 non-root aggregates and NAS data LIFs from node2 to the upgraded node1 and verify that the SAN LIFs exist on node1.

Steps

1. [Cable node1 for shared cluster-HA storage \(AFF A800 upgrade only\)](#)
2. [Boot node1 with the replacement system modules](#)
3. [Verify the node1 installation](#)
4. [Restore key-manager configuration on the upgraded node1](#)
5. [Move node1 non-root aggregates and NAS data LIFs from node2 to the upgraded node1](#)

Cable node1 for shared cluster-HA and storage

If you are performing one of the following upgrades, you must connect the cluster, HA, storage, data, and management connections that were previously connected to the node1 on the existing system to the newly installed node1 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	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.

If the...	Then...
e0M and BMC IP addresses are on different subnets	<ol style="list-style-type: none"> 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	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	Do not connect	Do not connect
HA	e1b	Do not connect	Do not connect
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	AFF A250, AFF C250	AFF A30, AFF C30, AFF C60	AFF A50
Cluster and HA	e0c	e4a	e4a
Cluster and HA	e0d	e2a (Use e4b if there is no e2a)	e2a (Use e4b if there is no e2a)
Ethernet storage Ports	Any available port	e3a, e3b	e3a, e3b
SAS storage ports	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.

Boot node1 with the replacement system modules

Node1 with the replacement modules is now ready to boot. The supported replacement modules are listed in the [supported systems matrix](#).



When replacing controller modules, move all connections from the old to the replacement controller module.

When replacing the controller and NVRAM modules, move only the console and management connections.

Steps

1. (AFF A250, AFF C250, AFF A800, or AFF C800 upgrade only) At the LOADER prompt, enter maintenance mode:

```
boot_ontap maint
```

- a. Answer `y` to the mixed platform confirmation prompt.
- b. Answer `yes` to the confirmation prompt.
- c. Show the state of the 100GbE interfaces:

```
storage port show.
```

All 100GbE ports connected to NS224 shelves or storage switches should report as `storage` ports, as shown in the example output below.

```
*> storage port show
Port Type Mode      Speed (Gb/s) State   Status  VLAN ID
----
e8a  ENET storage 100 Gb/s   enabled online  30
e8b  ENET storage 100 Gb/s   enabled online  30
e11a ENET storage 100 Gb/s   enabled online  30
e11b ENET storage 100 Gb/s   enabled online  30
```

- d. Exit maintenance mode:

```
halt
```

2. 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](#).

3. Boot the node into the boot menu:

```
boot_ontap menu
```

4. Reassign the old node1 disks to the replacement node1 by entering "22/7" and selecting the hidden option `boot_after_controller_replacement` when the node stops at the boot menu.

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 2](#).
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>

Changing sysid of node nodel disks.

Fetches sanown old_owner_sysid = 536940063 and calculated old sys id


```
= 536940063
Partner sysid = 4294967295, owner sysid = 536940063
.
.
<output truncated>
.
.
varfs_backup_restore: restore using /mroot/etc/varfs.tgz
varfs_backup_restore: attempting to restore /var/kmip to the boot
device
varfs_backup_restore: failed to restore /var/kmip to the boot device
varfs_backup_restore: attempting to restore env file to the boot
device
varfs_backup_restore: successfully restored env file to the boot
device wrote key file "/tmp/rndc.key"
varfs_backup_restore: timeout waiting for login
varfs_backup_restore: Rebooting to load the new varfs
Terminated
<node reboots>

System rebooting...

.
.
Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...
Rebooting to load the restored env file...

.
System rebooting...

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



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 node1 installation

After you boot node1 with the replacement controller module, verify that it is installed correctly.

For AFF A250, AFF C250, AFF A800, or AFF C800 upgrades only, you map the physical ports from the existing node1 to the replacement node1 because the physical ports are changing between the existing and replacement controllers. The physical ports don't change for the other systems supported by this upgrade procedure.

About this task

You must wait for node1 to join quorum and then resume the controller replacement operation.

At this point in the procedure, the controller upgrade operation should have paused as node1 attempts to join quorum automatically.

Steps

1. Verify that node1 has joined quorum:

```
cluster show -node node1 -fields health
```

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

2. Verify that node1 is part of the same cluster as node2 and that it is healthy:

```
cluster show
```



If node1 has not joined quorum after you boot, wait five minutes and check again. Depending on the cluster connection, it might take some time for the port reachability scan to complete and move LIFs to their respective home ports.

If node1 is still not in quorum after five minutes, consider modifying the cluster port of the new node by placing it in "Cluster ipspace" using the diagnostic privilege command `network port modify <port_name> -ipspace Cluster`.

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 that it was in before node1 was halted to perform the physical tasks of installing new controllers and moving cables:

```
system controller replace show
```

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

5. Resume the controller replacement operation:

```
system controller replace resume
```

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

```
Cluster::*> system controller replace show
Node          Status          Error-Action
-----
Node1         Paused-for-intervention  Follow the instructions given
in
Node2         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 node1*.

7. With the controller replacement in a paused state, proceed to [Restore network configuration on node1](#).

Restore network configuration on node1

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

```
network port show -node node1
```

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 node1:

```
network port reachability show -node node1
```

You should see output like the following example:

```

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 examples, node1 booted after the controller replacement. The ports that display "no-reachability" have no physical connectivity. You must repair any ports with a reachability status other than ok.



During the upgrade, the network ports and their connectivity should not change. All ports should reside in the correct broadcast domains and the network port reachability should not change. However, before moving LIFs from node2 back to node1, you must verify the reachability and health status of the network ports.

- Repair the reachability for each of the ports on node1 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
```

- Physical ports
- VLAN ports

You should see output like the following example:

```
Cluster ::> reachability repair -node node1 -port elb
```

```
Warning: Repairing port "node1:elb" 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 the upgraded node1

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

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

2. Verify that the SVM-KEK key is restored to "true" on node1 before you relocate the data aggregates:

```
::> security key-manager key query -node node1 -fields restored -key  
-type SVM-KEK
```

Example

```
::> security key-manager key query -node node1 -fields restored -key  
-type SVM-KEK
```

node	vserver	key-server	key-id
restored			
-----	-----	-----	-----
node1	svm1	""	0000000000000000020000000000a008a81976
true			2190178f9350e071fbb90f000000000000000

Move node1 non-root aggregates and NAS data LIFs from node2 to the upgraded node1

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

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. After you bring node1 online, you must verify that the LIFs are healthy and located on the appropriate ports.

Steps

1. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

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

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

2. Perform a network reachability check:

```
network port reachability show -node node1
```

Confirm that all connected ports, including the interface group and VLAN ports, show their status as OK.

3. For the following upgrades, you must reassign the FCP SAN LIFs.

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

For all other system upgrades, proceed to [Step 4](#).

- a. Reassign FCP SAN LIFs used for FCP or FC-NVMe data access to the correct home ports:

```
network interface show -vserver <vserver_hosting_fcp_lifs>
```

- b. For LIFs with the current node as the upgraded node1 and the current port reports "status oper" as "-" (because the port existed on the AFF A800 node but does not exist on the AFF A90 node), modify the current port before it can be brought online.

Verify that physical connectivity is established to the FC target port where the FC LIF needs to be moved:

- i. Set the LIF status to "down":


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

ii. Modify the home port of the LIF:

```
network interface modify -vserver <vserver_name> -lif <lif_name> - home-
node <node1> -home-port <FC_target_port>
```

iii. Set the LIF status to "up":

```
network interface modify -vserver <vserver> -lif <lif_name> -status-admin
up
```

Repeat Substeps a and b for each FC SAN LIF that is home on node1.

4. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following checks:

- Cluster health check
- Cluster LIF status check

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

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

5. 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.

6. If necessary, restore and revert any displaced LIFs. List any displaced LIFs:

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

If any LIFs are displaced, restore the home node back to node1:

```
cluster controller-replacement network displaced-interface restore-home-node
```

7. 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 © 2025 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.