



## **Install and boot node3**

### **AFF and FAS Controller Upgrade**

NetApp  
July 23, 2021

# Table of Contents

Install and boot node3 ..... 1

# Install and boot node3

You need to install node3 in the rack, transfer node1's connections to node3, boot node3, and install ONTAP. You then need to reassign any of node1's spare disks, any disks belonging to the root volume, and any non-root aggregates not relocated to node2 earlier.

## About this task

You need to netboot node3 if it does not have the same version of ONTAP 9 that is installed on node1. After you install node3, 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. See [Prepare for netboot](#).

However, you do not need to netboot node3 if it has the same or a later version of ONTAP 9 that is installed on node1.



**Important:** If you are upgrading a V-Series system connected to storage arrays or a system with FlexArray Virtualization software that is connected to storage arrays, you need to complete [Step 1](#) through [Step 5](#), leave this section at [Step 6](#) and follow instructions in [Configure FC ports on node3](#) and [Check and configure UTA/UTA2 ports on node3](#) as needed, entering commands in maintenance mode. You must then return to this section and resume with [Step 7](#). However, if you are upgrading a system with storage disks, you need to complete this entire section and then go to [Configure FC ports on node3](#) and [Check and configure UTA/UTA2 ports on node3](#), entering commands at the cluster prompt.

## Steps

1. Make sure that you have rack space for node3.

If node1 and node2 were in separate chassis, you can put node3 in the same rack location as node1. However, if node1 was in the same chassis with node2, then you need to put node3 into its own rack space, preferably close to the location of node1.

2. Install node3 in the rack, following the *Installation and Setup Instructions* for your node model.



If you are upgrading to a system with both nodes in the same chassis, install node4 in the chassis as well as node3. If you do not, when you boot node3, the node will behave as if it were in a dual-chassis configuration, and when you boot node4, the interconnect between the nodes will not come up.

3. Cable node3, moving the connections from node1 to node3.

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

- *Installation and Setup Instructions* or *FlexArray Virtualization Installation Requirements and Reference* for the node3 platform.
- The appropriate disk shelf guide.
- The *ONTAP 9 High-Availability Configuration Guide*

Cable the following connections:

- Console (remote management port)
- Cluster ports

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



You might not need to move the interconnect card or the cluster interconnect cable connection from node1 to node3 because most platform models have a unique interconnect card model. For the MetroCluster configuration, you need to move the FC-VI cable connections from node1 to node3. If the new host does not have an FC-VI card, you might need to move the FC-VI card.

4. Turn on the power to node3, and then interrupt the boot process by pressing Ctrl-C at the console terminal to access the boot environment prompt.

If you are upgrading to a system with both nodes in the same chassis, node4 also reboots. However, you can disregard the node4 boot until later.



When you boot node3, you might see the following warning message:

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

5. If you see the warning message in [Step 4](#), take the following actions:
  - a. Check for any console messages that might indicate a problem other than a low NVRAM battery, and, if necessary, take any required corrective action.
  - b. Allow the battery to charge and the boot process to complete.



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

6. Take one of the following actions:

If your system...	Then...
Has disks and no back-end storage	Skip Step 7 through Step 12 and go to <a href="#">Step 13</a> .

If your system...	Then...
Is a V-Series system or a system with FlexArray Virtualization software connected to storage arrays	<p>a. Go to <a href="#">Set the FC or UTA/UTA2 configuration on node3</a> and complete the subsections <a href="#">Configure FC ports on node3</a> and <a href="#">Check and configure UTA/UTA2 ports on node3</a>, as appropriate to your system.</p> <p>b. Return to this section and complete the remaining steps, beginning with <a href="#">Step 7</a>.</p> <p><b>Important:</b> You must reconfigure FC onboard ports, CNA onboard ports, and CNA cards before you boot ONTAP on the V-Series or system with FlexArray Virtualization software.</p>

7. Add the FC initiator ports of the new node to the switch zones.

If your system has a tape SAN, then you need zoning for the initiators. See your storage array and zoning documentation for instructions.

8. Add the FC initiator ports to the storage array as new hosts, mapping the array LUNs to the new hosts.

See your storage array and zoning documentation for instructions.

9. Modify the World Wide Port Name (WWPN) values in the host or volume groups associated with array LUNs on the storage array.

Installing a new controller module changes the WWPN values associated with each onboard FC port.

10. If your configuration uses switch-based zoning, adjust the zoning to reflect the new WWPN values.
11. Verify that the array LUNs are now visible to node3:

```
sysconfig -v
```

The system displays all the array LUNs visible to each of the FC initiator ports. If the array LUNs are not visible, you will not be able to reassign disks from node1 to node3 later in this section.

12. Press Ctrl-C to display the boot menu and select maintenance mode.
13. At the Maintenance mode prompt, enter the following command:

```
halt
```

The system stops at the boot environment prompt.

14. Take one of the following actions:

If the system you are upgrading to is in a...	Then...
Dual-chassis configuration (with controllers in different chassis)	Go to <a href="#">Step 15</a> .

If the system you are upgrading to is in a...	Then...
Single-chassis configuration (with controllers in the same chassis)	<p>a. Switch the console cable from node3 to node4.</p> <p>b. 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.</p> <p>The power should already be on if both controllers are in the same chassis.</p> <p><b>Note:</b> Leave node4 at the boot environment prompt; you will return to node4 in <a href="#">Install and boot node4</a>.</p> <p>c. If you see the warning message displayed in <a href="#">Step 4</a>, follow the instructions in <a href="#">Step 5</a></p> <p>d. Switch the console cable back from node4 to node3.</p> <p>e. Go to <a href="#">Step 15</a>.</p>

15. Configure node3 for ONTAP:

```
set-defaults
```

16. If NetApp Storage Encryption (NSE) is in use on this configuration, the `setenv bootarg.storageencryption.support` command must be set to `true`, and the `kmip.init.maxwait` variable needs to be set to `off` to avoid a boot loop after the node1 configuration is loaded:

```
setenv bootarg.storageencryption.support true
```

```
setenv kmip.init.maxwait off
```

17. If the version of ONTAP installed on node3 is the same or later than the version of ONTAP 9 installed on node1, list and reassign disks to the new node3:

```
boot_ontap
```



**Warning:** If this new node has ever been used in any other cluster or HA pair, you must run `wipeconfig` before proceeding. Failure to do so might result in service outages or data loss. Contact technical support if the replacement controller was previously used, especially if the controllers were running ONTAP running in 7-Mode.

18. Press CTRL-C to display the boot menu.

19. Take one of the following actions:

If the system you are upgrading...	Then...
Does <i>not</i> have the correct or current ONTAP version on node3	Go to <a href="#">Step 20</a> .

If the system you are upgrading...	Then...
Has the correct or current version of ONTAP on node3	Go to <a href="#">Step 25</a> .

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



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

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

21. Perform netboot on node3:

For...	Then...
FAS/AFF8000 series systems	<pre>netboot http://&lt;web_server_ip&gt;/&lt;path_to_webaccessible_directory&gt;/netboot/kernel</pre>
All other systems	<pre>netboot http://&lt;web_server_ip&gt;/&lt;path_to_webaccessible_directory&gt;/&lt;ontap_version&gt;_image.tgz</pre>

The <path\_to\_the\_web-accessible\_directory> should lead to where you downloaded the <ontap\_version>\_image.tgz in [Step 1](#) in the section *Prepare for netboot*.



Do not interrupt the boot.

- 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 wrong image might install. This issue applies to all releases of ONTAP.

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

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

- Complete the following substeps:

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

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

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

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

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

- Select **(5) Maintenance mode boot** by entering `5`, and then enter `y` when prompted to continue with the boot.
- Before continuing, go to [Set the FC or UTA/UTA2 configuration on node3](#) to make any necessary changes to the FC or UTA/UTA2 ports on the node.

Make the changes recommended in those sections, reboot the node, and go into maintenance mode.

- Find the system ID of node3:

```
disk show -a
```



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

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



You might see the message `disk show: No disks match option -a.` after entering the command. This is not an error message so you can continue with the procedure.

28. Reassign node1's spares, any disks belonging to the root, and any non-root aggregates that were not relocated to node2 earlier in [Relocate non-root aggregates from node1 to node2](#).

Enter the appropriate form of the `disk reassign` command based on whether your system has shared disks:

If disk type is...	Then run the command...
With shared disks	<code>disk reassign -s &lt;node1_sysid&gt; -d &lt;node3_sysid&gt; -p &lt;node2_sysid&gt;</code>
Without shared disks	<code>disk reassign -s &lt;node1_sysid&gt; -d &lt;node3_sysid&gt;</code>

For the `<node1_sysid>` value, use the information captured in [Record node1 information](#). To obtain the value for `<node3_sysid>`, use the `sysconfig` command.



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

The `disk reassign` command reassigns only those disks for which `<node1_sysid>` is the current owner.

The system displays the following message:

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

29. Enter `n`.

The system displays the following message:

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

30. Enter `y`

The system displays the following message:

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

31. Enter `y`.

32. If you are upgrading from a system with external disks to a system that supports internal and external disks (AFF A800 systems, for example), set the node1 aggregate as root to ensure node3 boots from the root aggregate of node1.



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

The following procedure sets node3 to boot from the root aggregate of node1:

a. Check the RAID, plex, and checksum information for the node1 aggregate:

```
aggr status -r
```

b. Check the status of the node1 aggregate:

```
aggr status
```

c. Bring the node1 aggregate online, if necessary:

```
aggr_online <root_aggr_from_node1>
```

d. Prevent the node3 from booting from its original root aggregate:

```
aggr offline <root_aggr_on_node3>
```

- e. Set the node1 root aggregate as the new root aggregate for node3:

```
aggr options <aggr_from_node1> root
```

- f. Verify that the root aggregate of node3 is offline and the root aggregate for the disks brought over from node1 is online and set to root:

```
aggr status
```



Failing to perform the previous substep might cause node3 to boot from the internal root aggregate, or it might cause the system to assume a new cluster configuration exists or prompt you to identify one.

The following shows an example of the command output:

```
-----  
Aggr State           Status           Options  
aggr0_nst_fas8080_15 online    raid_dp, aggr   root, nosnap=on  
                    fast zeroed  
                    64-bit  
  
aggr0 offline       raid_dp, aggr   diskroot  
                    fast zeroed  
                    64-bit  
-----
```

33. 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:

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

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

If the controller and chassis are not configured as "ha", use the following commands to correct the configuration:

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

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

If you have a MetroCluster configuration, use the following commands to modify the controller and chassis:

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

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

34. Destroy the mailboxes on node3:

```
mailbox destroy local
```

The console displays the following message:

```
Destroying mailboxes forces a node to create new empty mailboxes, which
clears any takeover state, removes all knowledge of out-of-date plexes
of mirrored volumes, and will prevent management services from going
online in 2-node cluster HA configurations. Are you sure you want to
destroy the local mailboxes?
```

35. Enter `y` at the prompt to confirm that you want to destroy the local mailboxes.

36. Exit maintenance mode:

```
halt
```

The system stops at the boot environment prompt.

37. On node2, check the system date, time, and time zone:

```
date
```

38. On node3, check the date at the boot environment prompt:

```
show date
```

39. If necessary, set the date on node3:

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

40. On node3, check the time at the boot environment prompt:

```
show time
```

41. If necessary, set the time on node3:

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

42. Verify the partner system ID is set correctly as noted in [Step 28](#) under `-p` switch:

```
printenv partner-sysid
```

43. If necessary, set the partner system ID on node3:

```
setenv partner-sysid <node2_sysid>
```

Save the settings:

```
saveenv
```

44. Access the boot menu at the boot environment prompt:

```
boot_ontap menu
```

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

The system displays the following message:

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

46. Enter `y` at the prompt.

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



The system might reboot twice before displaying the mismatch warning.

47. Confirm the mismatch as shown in the following example:

```
WARNING: System id mismatch. This usually occurs when replacing CF or  
NVRAM cards!  
Override system id (y|n) ? [n] y
```

The node might go through one round of reboot before booting normally.

48. Log in to node3.

## Copyright Information

Copyright © 2021 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## 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.