



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 must install node3 in the rack, transfer node1's connections to node3, boot node3, and install ONTAP. You must then reassign any of node1's spare disks, any disks belonging to the root volume, and any non-root aggregates that were not relocated to node2 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 largely automated; the operation pauses to allow you to check its status. You must manually resume the operation. In addition, you must verify the SAN LIFs have successfully moved to node3.

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, by following the instructions in [Prepare for netboot](#).

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 21](#), then leave this section and follow instructions in the [Configure FC ports on node3](#) and [Check and configure UTA/UTA2 ports on node3](#) sections as needed, entering commands in Maintenance mode. You must then return to this section and resume with [Step 23](#).
- If you are upgrading a system with storage disks, you need to complete this entire section and then go to the [Configure FC ports on node3](#) and [Check and configure UTA/UTA2 ports on node3](#) sections, 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.

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

Refer to [References](#) to link to the *FlexArray Virtualization Installation Requirements and Reference* and the *ONTAP 9 High-Availability Configuration Guide*.

- 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.



Refer to [Prepare for netboot](#).

6. 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 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: <pre>ifconfig e0M -auto</pre>
Not running	Manually configure the connection by using the following command at the boot environment prompt: <pre>ifconfig e0M -addr=<filer_addr> -mask=<netmask> -gw=<gateway> - dns=<dns_addr> domain=<dns_domain></pre> <p><filer_addr> is the IP address of the storage system. <netmask> is the network mask of the storage system. <gateway> is the gateway for the storage system. <dns_addr> is the IP address of a name server on your network. This parameter is optional. <dns_domain> is the Domain Name Service (DNS) domain name. This parameter is optional.</p> <p>Note: Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p>

7. Perform netboot on node3:

For...	Then...
FAS/AFF8000 series systems	<pre>netboot http://<web_server_ip/path_to_web-accessible_directory>/netboot/kernel</pre>
All other systems	<pre>netboot http://<web_server_ip/path_to_web-accessible_directory>/<ontap_version>_image.tgz</pre>

The <path_to_the_web-accessible_directory> should lead to where you downloaded the <ontap_version>_image.tgz in 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 wrong image might install. This issue applies to all ONTAP releases.

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. Enter `y` to reboot 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.

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
```

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
```

14. Exit maintenance mode:

```
halt
```

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

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

date

16. On node3, check the date by using the following command at the boot environment prompt:

```
show date
```

17. If necessary, set the date on node3:

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

18. On node3, check the time by using the following command at the boot environment prompt:

```
show time
```

19. If necessary, set the time on node3:

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

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

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

- a. Save the settings:

```
saveenv
```

21. On the new node, in boot loader, the `partner-sysid` parameter needs to be set. For node3, `partner-sysid` needs to be that of node2. Verify the `partner-sysid` for node3:

```
printenv partner-sysid
```

22. Take one of the following actions:

If your system...	Description
Has disks and no back-end storage	Go to Step 23
Is a V-Series system or a system with FlexArray Virtualization software connected to storage arrays	<p>a. Go to section Setting the FC or UTA/UTA2 configuration on node3 and complete the subsections in this section.</p> <p>b. Return to this section and complete the remaining steps, beginning with Step 23.</p> <p>Important: 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>

23. 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. If required, modify the onboard ports to initiator by referring to the [Configuring FC ports on node3](#). See your storage array and zoning documentation for further instructions on zoning.

24. 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.

25. Modify the worldwide 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.

26. If your configuration uses switch-based zoning, adjust the zoning to reflect the new WWPN values.

27. 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
```

28. Boot node into boot menu:

```
boot_ontap menu
```

If you do not have FC or UTA/UTA2 configuration, execute [Check and configure UTA/UTA2 ports on node 3, Step 15](#) so that node3 can recognize node1's disks.

29. For a MetroCluster configuration, V-Series systems and systems with FlexArray Virtualization software connected to storage arrays, you must set and configure the FC or UTA/UTA2 ports on node3 to detect the disks attached to the node.

To complete this task, go to section [Set the FC or UTA/UTA2 configuration on 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.