



Stage 2. Relocate resources and retire node1

Upgrade controllers

NetApp
February 22, 2024

Table of Contents

- Stage 2. Relocate resources and retire node1 1
 - Overview 1
 - Relocate non-root aggregates and NAS data LIFs owned by node1 to node2 1
 - Relocate failed or vetoed aggregates 2
 - Retire node1 3
 - Replace the node1 system modules 3
 - Netboot node1 10

Stage 2. Relocate resources and retire node1

Overview

During Stage 2, you relocate node1 non-root aggregates and NAS data LIFs to node2. This process is largely automated; the operation pauses to enable you to check its status. You must manually resume the operation. If required, you relocate failed or vetoed aggregates. You also record node1 information for use later in the procedure and then swap the corresponding node1 system modules, retire node1, and netboot the upgraded node1.

Steps

1. [Relocate non-root aggregates and NAS data LIFs owned by node1 to node2](#)
2. [Relocate failed or vetoed aggregates](#)
3. [Retire node1](#)
4. [Replace the node1 system modules](#)
5. [Netboot node1](#)

Relocate non-root aggregates and NAS data LIFs owned by node1 to node2

Before you can replace node1 with the replacement modules for your system upgrade, you must move the non-root aggregates and NAS data LIFs from node1 to node2 before eventually restoring the node1 resources back on node1 running on the replacement system. This process is largely automated; the operation pauses to enable you to check its status.

Before you begin

The operation should already be paused when you begin the task; you must manually resume the operation.

About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. You are not required to move SAN LIFs for cluster or service health during the upgrade. You must verify that the LIFs are healthy and located on appropriate ports after you bring node1 online as the replacement system.



The home owner for the aggregates and LIFs is not modified; only the current owner is modified.

Steps

1. Resume the aggregate relocation and NAS data LIF move operations:

```
system controller replace resume
```

All the non-root aggregates and NAS data LIFs are migrated from node1 to node2.

The operation pauses to enable you to verify whether all node1 non-root aggregates and non-SAN data LIFs have been migrated to node2.

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

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

3. With the operation still paused, verify that all the non-root aggregates are online for their state on node2:

```
storage aggregate show -node node2 -state online -root false
```

The following example shows that the non-root aggregates on node2 are online:

```
cluster::> storage aggregate show -node node2 state online -root false
```

Aggregate	Size	Available	Used%	State	#Vols	Nodes	RAID	Status
aggr_1	744.9GB	744.8GB	0%	online	5	node2		
raid_dp,normal								
aggr_2	825.0GB	825.0GB	0%	online	1	node2		
raid_dp,normal								

2 entries were displayed.

If the aggregates have gone offline or become foreign on node2, bring them online by using the following command on node2, once for each aggregate:

```
storage aggregate online -aggregate aggr_name
```

4. Verify that all the volumes are online on node2 by using the following command on node2 and examining its output:

```
volume show -node node2 -state offline
```

If any volumes are offline on node2, bring them online by using the following command on node2, once for each volume:

```
volume online -vserver vservice_name -volume volume_name
```

The *vservice_name* to use with this command is found in the output of the previous `volume show` command.

5. If any LIFs are down, set the administrative status of the LIFs to `up` by using the following command, once for each LIF:

```
network interface modify -vserver vservice_name -lif LIF_name -home-node  
nodename -status-admin up
```

Relocate failed or vetoed aggregates

If any aggregates fail to relocate or are vetoed, you must manually relocate the aggregates, or if necessary, override either the vetoes or destination checks.

About this task

The relocation operation will have paused due to the error.

Steps

1. Check the event management system (EMS) logs to determine why the aggregate failed to relocate or was vetoed.
2. Relocate any failed or vetoed aggregates:

```
storage aggregate relocation start -node node1 -destination node2 -aggregate  
-list aggr_name -ndo-controller-upgrade true
```

3. When prompted, enter *y*.
4. You can force relocation by using one of the following methods:

Option	Description
Overriding veto checks	Use the following command: <pre>storage aggregate relocation start -node <i>node1</i> -destination <i>node2</i> -aggregate-list <i>aggr_list</i> -ndo -controller-upgrade true -override-vetoes true</pre>
Overriding destination checks	Use the following command: <pre>storage aggregate relocation start -node <i>node1</i> -destination <i>node2</i> -aggregate-list <i>aggr_list</i> -ndo -controller-upgrade true -override-vetoes true -override-destination-checks true</pre>

Retire node1

To retire node1, you resume the automated operation to disable the HA pair with node2 and shut down node1 correctly.

Steps

1. Resume the operation:

```
system controller replace resume
```

2. Verify that node1 has been halted:

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

After node1 has completely halted, node1 should be at the **LOADER>** prompt. To see the **LOADER>** prompt, connect to the serial console of node1.

Replace the node1 system modules

Replace the node1 system modules for your upgrade configuration:

- [Replace the AFF A220, AFF A200, AFF C190, FAS2620, or FAS2720 controller module](#)



You can also use this procedure to replace an AFF A220 configured as an ASA.

- [Replace the AFF A700 or FAS9000 controller and NVRAM modules](#)



You can also use this procedure to replace an AFF A700 configured as an ASA.

Replace the AFF A220, AFF A200, AFF C190, FAS2620, or FAS2720 controller module

At this stage, node1 is down and all data is served by node2. Because node1 and node2 are in the same chassis and powered by the same set of power supplies, do NOT power off the chassis. You must take care to remove only the node1 controller module. Typically, node1 is controller A, located on the left side of the chassis when looking at the controllers from the rear of the system. The controller label is located on the chassis directly above the controller module.

Before you begin

If you are not already grounded, correctly ground yourself.

Remove the AFF A220, AFF A200, AFF C190, FAS2620, or FAS2720 controller module


To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps


1. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

2. Remove and set aside the cable management devices from the left and right sides of the controller module.



3. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.



4. Turn the controller module over and place it on a flat, stable surface.

Install the ASA A150, AFF A150, or FAS2820 controller module

Use the following procedure to install the ASA A150, AFF A150, or FAS2820 controller module in node1.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.



Do not completely insert the controller module in the chassis until instructed to do so later in the procedure.

2. Cable the management and console ports to the node1 controller module.



Because the chassis is already powered ON, node1 starts BIOS initialization followed by autoboot as soon as it is fully seated. To interrupt the node1 boot, before completely inserting the controller module into the slot, it is recommended that you connect the serial console and management cables to the node1 controller module.

3. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated. The locking latch rises when the controller module is fully seated. Close the cam handle to the locked position.



To avoid damaging the connectors, do not use excessive force when sliding the controller module into the chassis.

4. Connect the serial console as soon as the module is seated and be ready to interrupt AUTOBOOT of node1.
5. After you interrupt AUTOBOOT, node1 stops at the LOADER prompt. If you do not interrupt AUTOBOOT on time and node1 starts booting, wait for the prompt and press **Ctrl-C** to go into the boot menu. After the node stops at the boot menu, use option 8 to reboot the node and interrupt the AUTOBOOT during reboot.
6. At the LOADER> prompt of node1, set the default environment variables:

```
set-defaults
```

7. Save the default environment variables settings:

Replace the AFF A700 or FAS9000 controller and NVRAM modules

At this stage, node1 is down and all data is served by node2. Because node1 and node2 are in the same chassis and powered by the same set of power supplies, do NOT power off the chassis. You must take care to remove only the node1 controller module and the node1 NVRAM module. Typically, node1 is controller A, located on the left side of the chassis when looking at the controllers from the rear of the system. The controller label is located on the chassis directly above the controller module.

Before you begin

If you are not already grounded, correctly ground yourself.

Remove the AFF A700 or FAS9000 controller module

Use the following procedure to remove the AFF A700 or FAS9000 controller module.


Steps

1. Detach the console cable, if any, and the management cable from the node1 controller module before removing the controller module from node1.



When you are working on node1, you only remove the console and e0M cables from node1. You must not remove or change any other cables or connections on either node1 or node2 during this process.

2. Unlock and remove the controller module A from the chassis.
 - a. Slide the orange button on the cam handle downward until it unlocks.



1

Cam handle release button

2	Cam handle
---	------------

- b. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Remove the AFF A700 or FAS9000 NVRAM module

Use the following procedure to remove the AFF A700 or FAS9000 NVRAM module.



The AFF A700 or FAS9000 NVRAM module is in slot 6 and is double the height of the other modules in the system.

Steps

1. Unlock and remove the NVRAM module from slot 6 of node1.


- a. Depress the lettered and numbered cam button.

The cam button moves away from the chassis.

- b. Rotate the cam latch down until it is in a horizontal position.

The NVRAM module disengages from the chassis and moves a few inches.

- c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module face.



1	Lettered and numbered I/O cam latch
2	I/O latch completely unlocked

Install the ASA A900, AFF A900, or FAS9500 NVRAM and controller modules

Install the ASA A900, AFF A900, or FAS9500 NVRAM and controller modules that you received for the upgrade on node1.

You must note the following when performing the installation:

- Move all blank filler modules in slots 6-1 and 6-2 from the old NVRAM module to the new NVRAM module.
- Do NOT move the coredump device from the AFF A700 NVRAM module to the ASA A900 or AFF A900 NVRAM module.
- Move all flash cache modules installed in the FAS9000 NVRAM module to the FAS9500 NVRAM module.

Before you begin


If you are not already grounded, correctly ground yourself.

Install the ASA A900, AFF A900, or FAS9500 NVRAM module

Use the following procedure to install the ASA A900, AFF A900, or FAS9500 NVRAM module in slot 6 of node1.

Steps

1. Align the NVRAM module with the edges of the chassis opening in slot 6.
2. Gently slide the NVRAM module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the NVRAM module in place.



1	Lettered and numbered I/O cam latch
2	I/O latch completely unlocked

Install the ASA A900, AFF A900, or FAS9500 controller module on node1.

Use the following procedure to install the ASA A900, AFA A900, or FAS9500 controller module in node1.

Steps

- 1. Align the end of the controller module with opening A in the chassis, and then gently push the controller module halfway into the system.



Do not completely insert the controller module in the chassis until instructed to do so later in the procedure.

- 2. Cable the management and console ports to the node1 controller module.




Because the chassis is already powered ON, node1 starts BIOS initialization followed by autoboot as soon as it is fully seated. To interrupt the node1 boot, before completely inserting the controller module into the slot, it is recommended that you connect the serial console and management cables to the node1 controller module.

- 3. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latch rises when the controller module is fully seated.



To avoid damaging the connectors, do not use excessive force when sliding the controller module into the chassis.



1	Cam handle locking latch
2	Cam handle in the unlocked position

- 4. Connect the serial console as soon as the module is seated and be ready to interrupt AUTOBOOT of

node1.

5. After you interrupt AUTOBOOT, node1 stops at the LOADER prompt. If you do not interrupt AUTOBOOT on time and node1 starts booting, wait for the prompt and press **Ctrl-C** to go into the boot menu. After the node stops at the boot menu, use option 8 to reboot the node and interrupt the AUTOBOOT during reboot.
6. At the LOADER> prompt of node1, set the default environment variables:

```
set-defaults
```

7. Save the default environment variables settings:

```
saveenv
```

Netboot node1

After swapping the corresponding replacement system modules, you must netboot node1. The term netboot means that you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you add a copy of the ONTAP 9 boot image onto a web server that the system can access.

It is not possible to check the version of ONTAP installed on the boot media of the replacement controller module unless it is installed in a chassis and powered ON. The ONTAP version on the replacement system boot media must be same as the ONTAP version running on the old system that you are upgrading and both the primary and backup boot images must match. To verify the minimum supported ONTAP version for your upgrade, see [Overview](#).

You can configure the images by performing a netboot followed by the `wipeconfig` command from the boot menu. If the controller module was previously used in another cluster, the `wipeconfig` command clears any residual configuration on the boot media.

You can also use the USB boot option to perform the netboot. See the Knowledge Base article [How to use the boot_recovery LOADER command for installing ONTAP for initial setup of a system](#).

Before you begin

- Verify that you can access a HTTP server with the system.
- Download the necessary system files for your system and the correct version of ONTAP from the *NetApp Support Site*. Refer to [References](#) to link to the *NetApp Support Site*.

About this task

You must netboot the new controllers if they do not have the same version of ONTAP 9 installed on them that is installed on the original controllers. After you install each new controller, you boot the system 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.


Steps

1. Refer to [References](#) to link to the *NetApp Support Site* to download the files used for performing the netboot of the system.
2. Download the appropriate ONTAP software from the software download section of the *NetApp Support Site* and store the `<ontap_version>_image.tgz` file on a web-accessible directory.
3. Change to the web-accessible directory and verify that the files you need are available.

4. Your directory listing should contain `<ontap_version>_image.tgz`.
5. 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>
Not running	<p>Manually configure the connection by using the following command at the boot environment prompt: <code>ifconfig e0M -addr=<i>filer_addr</i> -mask=<i>netmask</i> -gw=<i>gateway</i> -dns=<i>dns_addr</i> -domain=<i>dns_domain</i></code></p> <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 Domain Name Service (DNS) domain name (optional).</p> <div>  <p>Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p> </div>

6. Perform netboot on node1:

```
netboot http://<web_server_ip/path_to_web_accessible_directory>/netboot/kernel
```



Do not interrupt the boot.

7. Wait for the node1 now running on the ASA A900, AFF A900, or FAS9500 controller module to boot and display the boot menu options as shown below:

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)?

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

This note applies to nondisruptive ONTAP software upgrades, and not controller upgrades.



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
```

The `<path_to_the_web-accessible_directory>` should lead to where you downloaded the `<ontap_version>_image.tgz` in [Step 2](#).

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 must be restored.

11. At the prompt, run the `wipeconfig` command to clear any previous configuration on the boot media:

- a. When you see the message below, answer `yes`:

```
This will delete critical system configuration, including cluster
membership.
Warning: do not run this option on a HA node that has been taken
over.
Are you sure you want to continue?:
```

- b. The node reboots to finish the `wipeconfig` and then stops at the boot menu.

12. Select option 5 to go to maintenance mode from the boot menu. Answer `yes` to the prompts until the node stops at maintenance mode and the command prompt `*>`.

13. 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
```

14. 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
```

15. Verify the `ha-config` settings:

```
ha-config show
```

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

16. Halt node1:

```
halt
```

Node1 should stop at the LOADER prompt.

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

```
date
```

18. On node1, check the date by using the following command at the boot environment prompt:

```
show date
```

19. If necessary, set the date on node1:

```
set date mm/dd/yyyy
```



Set the corresponding UTC date on node1.

20. On node1, check the time by using the following command at the boot environment prompt:

```
show time
```

21. If necessary, set the time on node1:

```
set time hh:mm:ss
```



Set the corresponding UTC time on node1.

22. Set the partner system ID on node1:

```
setenv partner-sysid node2_sysid
```

For node1, the `partner-sysid` must be that of node2. You can obtain the node2 system ID from the `node show -node node2` command output on node2.

- a. Save the settings:

```
saveenv
```

23. On node1, at the LOADER prompt, verify the `partner-sysid` for node1:

```
printenv partner-sysid
```


Copyright information

Copyright © 2024 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.