



Stage 2. Relocate and retire node1

Upgrade controllers

NetApp
August 29, 2024

Table of Contents

- Stage 2. Relocate and retire node1 1
 - Stage 2 overview 1
 - Relocate non-root aggregates and NAS data LIFs owned by node1 to node2 1
 - Relocate failed or vetoed aggregates 3
 - Retire node1 4
 - Prepare for netboot 4

Stage 2. Relocate and retire node1

Stage 2 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 the necessary node1 information, retire node1, and prepare to netboot node3 and node4 later in the procedure.

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. [Prepare for netboot](#)

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

Before you can replace node1 with node3, you must move the non-root aggregates and NAS data LIFs from node1 to node2 before eventually moving node1's resources to node3.

Before you begin

The operation must 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. Moving SAN LIFs is not necessary for cluster or service health during the upgrade. You must verify that the LIFs are healthy and located on appropriate ports after you bring node3 online.



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 vservers_name -volume volume_name
```

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

5. If the ports currently hosting data LIFs will not exist on the new hardware, remove them from the broadcast domain:

```
network port broadcast-domain remove-ports
```

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

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

7. If you have interface groups or VLANs configured, complete the following substeps:

- a. If you have not already saved them, record the VLAN and interface group information so you can re-create the VLANs and interface groups on node3 after node3 is booted up.

- b. Remove the VLANs from the interface groups:

```
network port vlan delete -node nodename -port ifgrp -vlan-id VLAN_ID
```



Follow the corrective action to resolve any errors that are suggested by the `vlan delete` command.

- c. Enter the following command and examine its output to see if there are any interface groups configured on the node:

```
network port ifgrp show -node nodename -ifgrp ifgrp_name -instance
```

The system displays interface group information for the node as shown in the following example:

```
cluster::> network port ifgrp show -node node1 -ifgrp a0a -instance
Node: node1
Interface Group Name: a0a
Distribution Function: ip
Create Policy: multimode_lacp
MAC Address: 02:a0:98:17:dc:d4
Port Participation: partial
Network Ports: e2c, e2d
Up Ports: e2c
Down Ports: e2d
```

- d. If any interface groups are configured on the node, record the names of those groups and the ports assigned to them, and then delete the ports by entering the following command, once for each port:

```
network port ifgrp remove-port -node nodename -ifgrp ifgrp_name -port
netport
```

Relocate failed or vetoed aggregates

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

About this task

The relocation operation will have paused due to the error.

Steps

1. Check the 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 * -ndocontroller-upgrade true
```

3. When prompted, enter `y`.

4. You can force relocation by using one of the following methods:

Option	Description
Overriding veto checks	Enter the following: <code>storage aggregate relocation start -override -vetoes * -ndocontroller-upgrade true</code>
Overriding destination checks	Enter the following: <code>storage aggregate relocation start -overridedestination-checks * -ndo -controllerupgrade true</code>

Retire node1

To retire node1, you resume the automated operation to disable the HA pair with node2 and shut node1 down correctly. Later in the procedure, you remove node1 from the rack or chassis.

Steps

1. Resume the operation:

```
system controller replace resume
```

2. Verify that node1 has been halted:

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

After you finish

You can decommission node1 after the upgrade is completed. See [Decommission the old system](#).

Prepare for netboot

After you physically rack node3 and node4 later in the procedure, you might need to netboot them. The term "netboot" means you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you put a copy of the ONTAP 9 boot image onto a web server that the system can access.

Before you begin

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

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.

However, you do not need to netboot the controllers if the same version of ONTAP 9 is installed on them that is installed on the original controllers. If so, you can skip this section and proceed to [Stage 3 Installing and booting node3](#)

Steps

1. Access 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.

For...	Then...
FAS/AFF8000 series systems	<p>Extract the contents of the <code><ontap_version>_image.tgz</code> file to the target directory:</p> <pre>tar -zxvf <ontap_version>_image.tgz</pre> <div> If you are extracting the contents on Windows, use 7-Zip or WinRAR to extract the netboot image.</div> <p>Your directory listing should contain a netboot folder with a kernel file: <code>netboot/kernel</code></p>
All other systems	<p>Your directory listing should contain the following file: <code><ontap_version>_image.tgz</code></p> <div> You do not need to extract the contents of the <code><ontap_version>_image.tgz</code> file.</div>

You will use the information in the directories in [Stage 3](#).

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.