



Stage 4. Relocate and retire node

AFF and FAS Controller Upgrade

NetApp
May 23, 2022

Table of Contents

- Stage 4. Relocate and retire node 1
- Stage 4. Relocate and retire node2 1
- Relocate non-root aggregates and NAS data LIFs from node2 to node3 1
- Retire node2 3

Stage 4. Relocate and retire node

Stage 4. Relocate and retire node2

During Stage 4, all non-root aggregates are relocated from node2 to node3; the non-SAN data LIFs owned by node2 are moved to node3. You must record the necessary node2 information and then retire node2.

Steps

1. [Relocate non-root aggregates and NAS data LIFs from node2 to node3](#)
2. [Retire node2](#)

Relocate non-root aggregates and NAS data LIFs from node2 to node3

Before you can replace node2 with node4, you must relocate the non-root aggregates that are owned by node2 to node3.

Before you begin

After the post-checks from the previous stage complete, the resource release for node2 starts automatically. The non-root aggregates and non-SAN data LIFs are migrated from node2 to node3.

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.

After the aggregates and LIFs are migrated, the operation is paused for verification purposes. At this stage, you must verify whether or not all the non-root aggregates and non-SAN data LIFs are migrated to node3.



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

Steps

1. Verify that all the non-root aggregates are online and their state on node3:

```
storage aggregate show -node <node3> -state online -root false
```

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

```
cluster::> storage aggregate show -node node3 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 node3, bring them online by using the following command on node3, once for each aggregate:

```
storage aggregate online -aggregate <aggr_name>
```

2. Verify that all the volumes are online on node3 by using the following command on node3 and examining the output:

```
volume show -node <node3> -state offline
```

If any volumes are offline on node3, bring them online by using the following command on node3, once for each volume:

```
volume online -vserver <vserver_name> -volume <volume_name>
```

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

3. Verify that the LIFs have been moved to the correct ports and have a status of up. 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 <vserver_name> - lif <LIF_name> -home-node <node_name> - status-admin up
```

4. 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
```

5. Verify that there are no data LIFs remaining on node2 by entering the following command and examining the output:

```
network interface show -curr-node <node2> -role data
```

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

- a. Record 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>
```

c. Check if there are any interface groups configured on the node by entering the following command and examining its output:

```
network port ifgrp show -node <node2> -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 node2 -ifgrp a0a -instance
                Node: node3
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>
```

Retire node2

To retire node2, you must shut node2 down correctly and remove it from the rack or chassis.

Steps

1. Resume the operation:

```
system controller replace resume
```

The node halts automatically.

After you finish

You can decommission node2 after the upgrade is completed. See [Decommission the old system](#).

Copyright Information

Copyright © 2022 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.