



Stage 4. Record information and retire node2

Upgrade controllers

NetApp
December 01, 2022

Table of Contents

- Stage 4. Record information and retire node2 1
- Stage 4. Record node2 information and retire node2 1
- Record node2 information 1
- Retire node2 4

Stage 4. Record information and retire node2

Stage 4. Record node2 information and retire node2

During Stage 4, you record node2 information and then retire node2.

Steps

1. [Record node2 information](#)
2. [Retire node2](#)

Record node2 information

Before you can shut down and retire node2, you must record information about its cluster network, management, and FC ports as well as its NVRAM System ID. You need that information later in the procedure when you map node2 to node4 and reassign disks.

Steps

1. Find the cluster network, node-management, intercluster, and cluster-management ports on node2:

```
network interface show -curr-node node_name -role  
cluster,intercluster,nodemgmt,cluster-mgmt
```

The system displays the LIFs for that node and other nodes in the cluster, as shown in the following example:

```
cluster::> network interface show -curr-node node2 -role  
cluster,intercluster,node-mgmt,cluster-mgmt
```

Is	Logical	Status	Network	Current	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
node2	intercluster	up/up	192.168.1.202/24	node2	e0e
true	clus1	up/up	169.254.xx.xx/24	node2	e0a
true	clus2	up/up	169.254.xx.xx/24	node2	e0b
true	mgmt1	up/up	192.168.0.xxx/24	node2	e0c
true					

4 entries were displayed.



Your system might not have intercluster LIFs. You will have a cluster management LIF only on one node of a node pair. A cluster management LIF was displayed in the example output of [Step 1](#) in *Record node1 port information*.

2. Capture the information in the output to use in the section [Map ports from node2 to node4](#).

The output information is required to map the new controller ports to the old controller ports.

3. Determine physical ports on node2:

```
network port show -node node_name -type physical +
```

node_name is the node which is being migrated.

The system displays the physical ports on node2, as shown in the following example:

```
cluster::> network port show -node node2 -type physical
```

(Mbps)	Port	IPspace	Broadcast Domain	Link	MTU	Speed
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper
node2						
	e0M	Default	IP_address	up	1500	auto/100
	e0a	Default	-	up	1500	auto/1000
	e0b	Default	-	up	1500	auto/1000
	e1a	Cluster	Cluster	up	9000	auto/10000
	e1b	Cluster	Cluster	up	9000	auto/10000

5 entries were displayed.

4. Record the ports and their broadcast domains.

The broadcast domains will need to be mapped to the ports on the new controller later in the procedure.

5. Determine the FC ports on node2:

```
network fcp adapter show
```

The system displays the FC ports on the node2, as shown in the following example:

```

cluster::> network fcp adapter show -node node2
                Connection  Host
Node           Adapter  Established  Port  Address
-----
node2
                0a      ptp         11400
node2
                0c      ptp         11700
node2
                6a      loop        0
node2
                6b      loop        0
4 entries were displayed.

```

6. Record the ports.

The output information is required to map the new FC ports on the new controller later in the procedure.

7. If you have not done so earlier, check whether there are interface groups or VLANs configured on node2:

```

ifgrp show

vlan show

```

You will use the information in the section [Map ports from node2 to node4](#).

8. Take one of the following actions:

If you...	Then...
Recorded NVRAM System ID number in Prepare the nodes for upgrade	Go to Retire node2 .
Did not record the NVRAM System ID number in Prepare the nodes for upgrade	Complete Step 9 and Step 10 and then go to the next section, Retire node2 .

9. Display the attributes of node2:

```

system node show -instance -node node2

```

```

cluster::> system node show -instance -node node2
                ...
                NVRAM System ID: system_ID
                ...

```

10. Record the NVRAM System ID to use in the section [Install and boot node4](#).

Retire node2

To retire node2, you must shut node2 down correctly and remove it from the rack or chassis. If the cluster is in a SAN environment, you also must delete the SAN LIFs.

Steps

1. Take one of the following actions:

If the cluster is...	Then...
A two-node cluster	Go to Step 2 .
A cluster with more than two nodes	Go to Step 9 .

2. Access the advanced privilege level by entering the following command on either node:

```
set -privilege advanced
```

3. Verify that the cluster HA has been disabled by entering the following command and examining its output:

```
cluster ha show
```

The system displays the following message:

```
High Availability Configured: false
```

4. Check if node2 currently holds epsilon by entering the following command and examining its output:

```
cluster show
```

The following example shows that node2 holds epsilon:

```
cluster*::> cluster show
Node           Health  Eligibility  Epsilon
-----
node1          true    true         false
node2          true    true         true
```

```
Warning: Cluster HA has not been configured. Cluster HA must be
configured on a two-node cluster to ensure data access availability in
the event of storage failover. Use the "cluster ha modify -configured
true" command to configure cluster HA.
```

```
2 entries were displayed.
```



If you are upgrading a HA pair in a cluster with multiple HA pairs, you must move epsilon to the node of a HA pair not undergoing a controller upgrade. For example, if you are upgrading nodeA/nodeB in a cluster with the HA pair configuration nodeA/nodeB and nodeC/nodeD, you must move epsilon to nodeC or nodeD.

5. If node2 holds epsilon, mark epsilon as `false` on the node so that it can be transferred to node3:

```
cluster modify -node node2 -epsilon false
```

6. Transfer epsilon to node3 by marking epsilon `true` on node3:

```
cluster modify -node node3 -epsilon true
```

7. Verify if the setup is a two-node switchless cluster:

```
network options switchless-cluster show
```

```
cluster::*> network options switchless-cluster show
Enable Switchless Cluster: false/true
```

The value of this command must match the physical state of the system.

8. Verify if the setup is a two-node switchless cluster:

```
network options switchless-cluster show
```

```
cluster::*> network options switchless-cluster show
Enable Switchless Cluster: false/true
```

The value of this command must match the physical state of the system.

9. Return to the admin level:

```
set -privilege admin
```

10. Halt node2 by entering the following command on either controller:

```
system node halt -node node2
```

11. After node2 shuts down completely, remove it from the chassis or the rack. 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.

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.