



Prepare the nodes for upgrade

AFF and FAS Controller Upgrade

NetApp
July 23, 2021

Table of Contents

Prepare the nodes for upgrade 1

Prepare the nodes for upgrade

Before you can replace the original nodes, you must ensure that they are in an HA pair, have no missing or failed disks, can access each other's storage, and do not own data LIFs assigned to the other nodes in the cluster. You also need to collect information about the original nodes and, if the cluster is in a SAN environment, ensure that all the nodes in the cluster are in quorum.

Steps

1. Ensure that each of the original nodes has enough resources to adequately support the workload of both nodes during takeover mode.

Refer to [References](#) to link to *ONTAP 9 High-Availability Configuration Guide* and follow the *Best practices for HA pairs* section. Neither of the original nodes should be running at more than 50 percent utilization; if a node is running at less than 50 percent utilization, it can handle the loads for both nodes during the controller upgrade.

2. Complete the following substeps to create a performance baseline for the original nodes:

- a. Make sure that the diagnostic user account is unlocked.

Important: The diagnostic user account is intended only for low-level diagnostic purposes and should be used only with guidance from technical support.

Important: For information about unlocking the user accounts, refer to [References](#) to link to the *System Administration Reference*.

- b. Refer to [References](#) to link to the *NetApp Support Site* and download the Performance and Statistics Collector (Perfstat Converged).

The Perfstat Converged tool lets you establish a performance baseline for comparison after the upgrade.

- c. Create a performance baseline, following the instructions on the NetApp Support Site.

3. Refer to [References](#) to link to the *NetApp Support Site* and open a support case on the NetApp Support Site.

You can use the case to report any issues that might arise during the upgrade.

4. Verify that NVMEM or NVRAM batteries of node3 and node4 are charged, and charge them if they are not.

You need to physically check node3 and node4 to see if the NVMEM or NVRAM batteries are charged. For information about the LEDs for the model of node3 and node4, refer to [References](#) to link to the *Hardware Universe*.



Attention Do not try to clear the NVRAM contents. If there is a need to clear the contents of NVRAM, contact NetApp technical support.

5. Check the version of ONTAP on node3 and node4.

The new nodes must have the same version of ONTAP 9.x installed on them that is installed on the original nodes. If the new nodes have a different version of ONTAP installed, you need to netboot the new controllers after you install them. For instructions on how to upgrade ONTAP, refer to [References](#) to link to

the *ONTAP 9 Upgrade and Revert/Downgrade Guide*.

Information about the version of ONTAP on node3 and node4 should be included in the shipping boxes. The ONTAP version is displayed when the node boots up or you can boot the node to maintenance mode and run the command:

```
version
```

6. Check whether you have two or four cluster LIFs on node1 and node 2:

```
network interface show -role cluster
```

The system displays any cluster LIFs, as shown in the following example:

```
cluster::> network interface show -role cluster
      Logical      Status      Network      Current      Current      Is
Vserver Interface  Admin/Oper  Address/Mask  Node         Port         Home
-----
node1
      clus1        up/up       172.17.177.2/24  node1        e0c          true
      clus2        up/up       172.17.177.6/24  node1        e0e          true
node2
      clus1        up/up       172.17.177.3/24  node2        e0c          true
      clus2        up/up       172.17.177.7/24  node2        e0e          true
```

7. If you have two or four cluster LIFs on node1 or node2, make sure that you can ping both cluster LIFs across all the available paths by completing the following substeps:

- a. Enter the advanced privilege level:

```
set -privilege advanced
```

The system displays the following message:

```
Warning: These advanced commands are potentially dangerous; use them
only when directed to do so by NetApp personnel.
Do you wish to continue? (y or n):
```

- b. Enter `y`.
- c. Ping the nodes and test the connectivity:

```
cluster ping-cluster -node node_name
```

The system displays a message similar to the following example:

```

cluster::*> cluster ping-cluster -node node1
Host is node1
Getting addresses from network interface table...
Local = 10.254.231.102 10.254.91.42
Remote = 10.254.42.25 10.254.16.228
Ping status:
...
Basic connectivity succeeds on 4 path(s) Basic connectivity fails on
0 path(s)
.....
Detected 1500 byte MTU on 4 path(s):
Local 10.254.231.102 to Remote 10.254.16.228
Local 10.254.231.102 to Remote 10.254.42.25
Local 10.254.91.42 to Remote 10.254.16.228
Local 10.254.91.42 to Remote 10.254.42.25
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

If the node uses two cluster ports, you should see that it is able to communicate on four paths, as shown in the example.

d. Return to the administrative level privilege:

```
set -privilege admin
```

8. Ensure that node1 and node2 are in an HA pair and verify that the nodes are connected to each other, and that takeover is possible:

```
storage failover show
```

The following example shows the output when the nodes are connected to each other and takeover is possible:

```

cluster:::> storage failover show

```

Node	Partner	Takeover Possible	State Description
node1	node2	true	Connected to node2
node2	node1	true	Connected to node1

Neither node should be in partial giveback. The following example shows that node1 is in partial giveback:

```

cluster::> storage failover show

Node           Partner           Takeover
Possible State Description
-----
node1          node2              true        Connected to node2, Partial
giveback
node2          node1              true        Connected to node1

```

If either node is in partial giveback, use the `storage failover giveback` command to perform the giveback, and then use the `storage failover show-giveback` command to make sure that no aggregates still need to be given back. For detailed information about the commands, refer to [References](#) to link to the *ONTAP 9 High-Availability Configuration Guide*.

9. Ensure that neither node1 nor node2 owns the aggregates for which it is the current owner (but not the home owner):

```

storage aggregate show -node <node_name> -is-home false -fields owner-
name,homeName,state

```

If neither node1 nor node2 owns aggregates for which it is the current owner (but not the home owner), the system will return a message similar to the following example:

```

cluster::> storage aggregate show -node node2 -is-home false -fields
owner-name,homeName,state
There are no entries matching your query.

```

The following example shows the output of the command for a node named node2 that is the home owner, but not the current owner, of four aggregates:

```

cluster::> storage aggregate show -node node2 -is-home false
-fields owner-name,home-name,state

aggregate    home-name    owner-name    state
-----
aggr1        node1        node2         online
aggr2        node1        node2         online
aggr3        node1        node2         online
aggr4        node1        node2         online

4 entries were displayed.

```

10. Take one of the following actions:

If the command in Step 9 ...	Then...
Had blank output	Skip Step 11 and go to Step 12 .

If the command in Step 9...	Then...
Had output	Go to Step 11 .

11. If either node1 or node2 owns aggregates for which it is the current owner but not the home owner, complete the following substeps:

- a. Return the aggregates currently owned by the partner node to the home owner node:

```
storage failover giveback -ofnode <home_node_name>
```

- b. Verify that neither node1 nor node2 still owns aggregates for which it is the current owner (but not the home owner):

```
storage aggregate show -nodes <node_name> -is-home false -fields owner-name,home-name,state
```

The following example shows the output of the command when a node is both the current owner and home owner of aggregates:

```
cluster::> storage aggregate show -nodes node1
           -is-home true -fields owner-name,home-name,state

aggregate      home-name      owner-name      state
-----
aggr1          node1          node1           online
aggr2          node1          node1           online
aggr3          node1          node1           online
aggr4          node1          node1           online

4 entries were displayed.
```

12. Ensure that node1 and node2 can access each other's storage and verify that no disks are missing:

```
storage failover show -fields local-missing-disks,partner-missing-disks
```

The following example shows the output when no disks are missing:

```
cluster::> storage failover show -fields local-missing-disks,partner-missing-disks

node      local-missing-disks  partner-missing-disks
-----
node1     None                None
node2     None                None
```

If any disks are missing, refer to [References](#) to link to the *ONTAP 9 Disks and Aggregates Power Guide*, the *ONTAP 9 Logical Storage Management Guide*, and the *ONTAP 9 High-Availability Configuration Guide*

to configure storage for the HA pair.

13. Ensure that node1 and node2 are healthy and eligible to participate in the cluster:

```
cluster show
```

The following example shows the output when both nodes are eligible and healthy:

```
cluster::> cluster show
```

Node	Health	Eligibility
node1	true	true
node2	true	true

14. Set the privilege level to advanced:

```
set -privilege advanced
```

15. Ensure that node1 and node2 are running the same ONTAP release:

```
system node image show -node <node1,node2> -iscurrent true
```

The following example shows the output of the command:

```
cluster::*> system node image show -node node1,node2 -iscurrent true
```

Node	Image	Is Default	Is Current	Version	Install Date
node1	image1	true	true	9.1	2/7/2017 20:22:06
node2	image1	true	true	9.1	2/7/2017 20:20:48

```
2 entries were displayed.
```

16. Verify that neither node1 nor node2 owns any data LIFs that belong to other nodes in the cluster and check the `Current Node` and `Is Home` columns in the output:

```
network interface show -role data -is-home false -curr-node <node_name>
```

The following example shows the output when node1 has no LIFs that are home-owned by other nodes in the cluster:


```
cluster::> network interface show -role data -is-home false -curr-node
node1
There are no entries matching your query.
```

The following example shows the output when node1 owns data LIFs home-owned by the other node:

```
cluster::> network interface show -role data -is-home false -curr-node
node1
```

Current Is	Logical	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
vs0	data1	up/up	172.18.103.137/24	node1	e0d
false	data2	up/up	172.18.103.143/24	node1	e0f
false					

2 entries were displayed.

- If the output in [Step 15](#) shows that either node1 or node2 owns any data LIFs home-owned by other nodes in the cluster, migrate the data LIFs away from node1 or node2:

```
network interface revert -vserver * -lif *
```

For detailed information about the `network interface revert` command, refer to [References](#) to link to the *ONTAP 9 Commands: Manual Page Reference*.

- Check whether node1 or node2 owns any failed disks:

```
storage disk show -nodelist <node1,node2> -broken
```

If any of the disks have failed, remove them, following instructions in the *ONTAP 9 Disks and Aggregates Power Guide*. (Refer to [References](#) to link to the *ONTAP 9 Disks and Aggregates Power Guide*.)

- Collect information about node1 and node2 by completing the following substeps and recording the output of each command:



You will use this information later in the procedure.

- Record the model, system ID, and serial number of both nodes:

```
system node show -node <node1,node2> -instance
```



You will use the information to reassign disks and decommission the original nodes.

- b. Enter the following command on both node1 and node2 and record information about the shelves, number of disks in each shelf, flash storage details, memory, NVRAM, and network cards from the output:

```
run -node <node_name> sysconfig
```



You can use the information to identify parts or accessories that you might want to transfer to node3 or node4. If you do not know if the nodes are V-Series systems or have FlexArray Virtualization software, you can learn that also from the output.

- c. Enter the following command on both node1 and node2 and record the aggregates that are online on both nodes:

```
storage aggregate show -node <node_name> -state online
```



You can use this information and the information in the following substep to verify that the aggregates and volumes remain online throughout the procedure, except for the brief period when they are offline during relocation.

- d. Enter the following command on both node1 and node2 and record the volumes that are offline on both nodes:

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



After the upgrade, you will run the command again and compare the output with the output in this step to see if any other volumes have gone offline.

20. Enter the following commands to see if any interface groups or VLANs are configured on node1 or node2:

```
network port ifgrp show
```

```
network port vlan show
```

Make note of whether interface groups or VLANs are configured on node1 or node2; you need that information in the next step and later in the procedure.

21. Complete the following substeps on both node1 and node2 to ensure that physical ports can be mapped correctly later in the procedure:

- a. Enter the following command to see if there are failover groups on the node other than `clusterwide`:

```
network interface failover-groups show
```

Failover groups are sets of network ports present on the system. Because upgrading the controller hardware can change the location of physical ports, failover groups can be inadvertently changed during the upgrade.

The system displays failover groups on the node, as shown in the following example:

```

cluster::> network interface failover-groups show

Vserver          Group          Targets
-----
Cluster          Cluster        node1:e0a, node1:e0b
                  node2:e0a, node2:e0b

fg_6210_e0c      Default        node1:e0c, node1:e0d
                  node1:e0e, node2:e0c
                  node2:e0d, node2:e0e

2 entries were displayed.

```

- b. If there are failover groups present other than `clusterwide`, record the failover group names and the ports that belong to the failover groups.
- c. Enter the following command to see if there are any VLANs configured on the node:

```
network port vlan show -node <node_name>
```

VLANs are configured over physical ports. If the physical ports change, then the VLANs will need to be re-created later in the procedure.

The system displays VLANs configured on the node, as shown in the following example:

```

cluster::> network port vlan show

Network Network
Node      VLAN Name Port      VLAN ID MAC Address
-----
node1     e1b-70   e1b       70       00:15:17:76:7b:69

```

- d. If there are VLANs configured on the node, take note of each network port and VLAN ID pairing.

22. Take one of the following actions:

If interface groups or VLANs are...	Then...
On node1 or node2	Complete Step 23 and Step 24 .
Not on node1 or node2	Go to Step 24 .

- 23. If you do not know if node1 and node2 are in a SAN or non-SAN environment, enter the following command and examine its output:

```
network interface show -vserver <vserver_name> -data-protocol iscsi|fc
```

If neither iSCSI nor FC is configured for the SVM, the command will display a message similar to the

following example:

```
cluster::> network interface show -vserver Vserver8970 -data-protocol
iscsi|fc
There are no entries matching your query.
```

You can confirm that the node is in a NAS environment by using the `network interface show` command with the `-data-protocol nfs|cifs` parameters.

If either iSCSI or FC is configured for the SVM, the command will display a message similar to the following example:

```
cluster::> network interface show -vserver vs1 -data-protocol iscsi|fc
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
vs1	vs1_lif1	up/down	172.17.176.20/24	node1	0d	true

24. Verify that all the nodes in the cluster are in quorum by completing the following substeps:

a. Enter the advanced privilege level:

```
set -privilege advanced
```

The system displays the following message:

```
Warning: These advanced commands are potentially dangerous; use them
only when directed to do so by NetApp personnel.
Do you wish to continue? (y or n):
```

b. Enter `y`.

c. Verify the cluster service state in the kernel, once for each node:

```
cluster kernel-service show
```

The system displays a message similar to the following example:

```

cluster::*> cluster kernel-service show

Master      Cluster      Quorum      Availability  Operational
Node        Node         Status      Status       Status
-----
node1       node1        in-quorum   true         operational
            node2        in-quorum   true         operational

2 entries were displayed.

```

Nodes in a cluster are in quorum when a simple majority of nodes are healthy and can communicate with each other. For more information, refer to [References](#) to link to the *System Administration Reference*.

d. Return to the administrative privilege level:

```
set -privilege admin
```

25. Take one of the following actions:

If the cluster...	Then...
Has SAN configured	Go to Step 26 .
Does not have SAN configured	Go to Step 29 .

26. Verify that there are SAN LIFs on node1 and node2 for each SVM that has either SAN iSCSI or FC service enabled by entering the following command and examining its output:

```
network interface show -data-protocol iscsi|fc -home-node <node_name>
```

The command displays SAN LIF information for node1 and node2. The following examples show the status in the Status Admin/Oper column as up/up, indicating that SAN iSCSI and FC service are enabled:

```

cluster::> network interface show -data-protocol iscsi|fc
          Logical      Status      Network      Current
Current Is
Vserver   Interface  Admin/Oper  Address/Mask      Node
Port      Home
-----
-----
a_vs_iscsi data1      up/up      10.228.32.190/21  node1      e0a
true
          data2      up/up      10.228.32.192/21  node2      e0a
true

b_vs_fcp   data1      up/up      20:09:00:a0:98:19:9f:b0  node1      0c
true
          data2      up/up      20:0a:00:a0:98:19:9f:b0  node2      0c
true

c_vs_iscsi_fcp data1      up/up      20:0d:00:a0:98:19:9f:b0  node2      0c
true
          data2      up/up      20:0e:00:a0:98:19:9f:b0  node2      0c
true
          data3      up/up      10.228.34.190/21  node2      e0b
true
          data4      up/up      10.228.34.192/21  node2      e0b
true

```

Alternatively, you can view more detailed LIF information by entering the following command:

```
network interface show -instance -data-protocol iscsi|fc
```

27. Capture the default configuration of any FC ports on the original nodes by entering the following command and recording the output for your systems:

```
ucadmin show
```

The command displays information about all FC ports in the cluster, as shown in the following example:

```
cluster::> uadmin show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
node1	0a	fc	initiator	-	-	online
node1	0b	fc	initiator	-	-	online
node1	0c	fc	initiator	-	-	online
node1	0d	fc	initiator	-	-	online
node2	0a	fc	initiator	-	-	online
node2	0b	fc	initiator	-	-	online
node2	0c	fc	initiator	-	-	online
node2	0d	fc	initiator	-	-	online

8 entries were displayed.

You can use the information after the upgrade to set the configuration of FC ports on the new nodes.

28. If you are upgrading a V-Series system or a system with FlexArray Virtualization software, capture information about the topology of the original nodes by entering the following command and recording the output:

```
storage array config show -switch
```

The system displays topology information, as show in the following example:

```
cluster::> storage array config show -switch
```

Node	Grp	LUN Cnt	Array Name	Array Target	Port	Switch	Port	Initiator
node1	0	50	I_1818FASTT_1	205700a0b84772da		vgbr6510a	5	
			vgbr6510s164:3	0d				
			vgbr6510s164:4	2b		vgbr6510a	6	
			vgbr6510s163:1	0c		vgbr6510b	6	
node2	0	50	I_1818FASTT_1	205700a0b84772da		vgbr6510a	5	
			vgbr6510s164:1	0d				
			vgbr6510s164:2	2b		vgbr6510a	6	
			vgbr6510s163:3	0c		vgbr6510b	6	
			vgbr6510s163:4	2a		vgbr6510b	5	

7 entries were displayed.

29. Complete the following substeps:

a. Enter the following command on one of the original nodes and record the output:

```
service-processor show -node * -instance
```

The system displays detailed information about the SP on both nodes.

b. Ensure that the SP status is `online`.

c. Ensure that the SP network is configured.

d. Record the IP address and other information about the SP.

You might want to reuse the network parameters of the remote management devices, in this case the SPs, from the original system for the SPs on the new nodes.

For detailed information about the SP, refer to [References](#) to link to the *System Administration Reference* and the *ONTAP 9 Commands: Manual Page Reference*.

30. If you want the new nodes to have the same licensed functionality as the original nodes, enter the following command to see the cluster licenses on the original system:

```
system license show -owner *
```


The following example shows the site licenses for cluster1:

```

system license show -owner *
Serial Number: 1-80-000013
Owner: cluster1

Package           Type      Description           Expiration
-----
Base              site     Cluster Base License -
NFS               site     NFS License           -
CIFS              site     CIFS License          -
SnapMirror        site     SnapMirror License    -
FlexClone         site     FlexClone License     -
SnapVault         site     SnapVault License     -
6 entries were displayed.

```

31. Obtain new license keys for the new nodes at the *NetApp Support Site*. Refer to [References](#) to link to *NetApp Support Site*.

If the site does not have the license keys you need, contact your NetApp sales representative.

32. Check whether the original system has AutoSupport enabled by entering the following command on each node and examining its output:

```
system node autosupport show -node <node1,node2>
```

The command output shows whether AutoSupport is enabled, as shown in the following example:

```

cluster::> system node autosupport show -node node1,node2

Node           State      From           To           Mail Hosts
-----
node1          enable    Postmaster     admin@netapp.com  mailhost
node2          enable    Postmaster     -            mailhost
2 entries were displayed.

```

33. Take one of the following actions:

If the original system...	Then...
Has AutoSupport enabled...	<ul style="list-style-type: none"> a. Go to Step 34. b. Go to the section Get an IP address of an external key management server for Storage Encryption.

If the original system...	Then...
Does not have AutoSupport enabled...	<p>a. Enable AutoSupport by following the instructions in the <i>System Administration Reference</i>. (Refer to References to link to the <i>System Administration Reference</i>.)</p> <p>Note: AutoSupport is enabled by default when you configure your storage system for the first time. Although you can disable AutoSupport at any time, you should leave it enabled. Enabling AutoSupport can significantly help identify problems and solutions should a problem occur on your storage system.</p> <p>b. Go to the Get an IP address of an external key management server for Storage Encryption section.</p>

34. Verify that AutoSupport is configured with the correct mailhost details and recipient e-mail IDs by entering the following command on both of the original nodes and examining the output:

```
system node autosupport show -node node_name -instance
```

For detailed information about AutoSupport, refer to [References](#) to link to the *System Administration Reference* and the *ONTAP 9 Commands: Manual Page Reference*.

35. Send an AutoSupport message to NetApp for node1 by entering the following command:

```
system node autosupport invoke -node node1 -type all -message "Upgrading node1 from platform_old to platform_new"
```



Do not send an AutoSupport message to NetApp for node2 at this point; you do so later in the procedure.

36. Verify that the AutoSupport message was sent by entering the following command and examining its output:

```
system node autosupport show -node <node1> -instance
```

The fields `Last Subject Sent:` and `Last Time Sent:` contain the message title of the last message sent and the time the message was sent.

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.