



Manual nondisruptive (rolling method) using the CLI

ONTAP 9

NetApp
May 06, 2021

Table of Contents

- Manual nondisruptive (rolling method) using the CLI 1
 - Updating the first node in an HA pair 1
 - Updating the partner node in an HA pair 6

Manual nondisruptive (rolling method) using the CLI

The rolling upgrade method enables you to update a cluster of two or more nodes nondisruptively. This method has several steps: initiating a failover operation on each node in an HA pair, updating the “failed” node, initiating giveback, and then repeating the process for each HA pair in the cluster.

You must have satisfied upgrade preparation requirements.

1. Update the first node in an HA pair

You upgrade the first node in an HA pair by initiating a takeover by the node’s partner. The partner serves the node’s data while the first node is upgraded.

2. Update the second node in an HA pair

After upgrading or downgrading the first node in an HA pair, you upgrade its partner by initiating a takeover on it. The first node serves the partner’s data while the partner node is upgraded.

3. Repeat these steps for each additional HA pair.

You should complete post-upgrade tasks.

Updating the first node in an HA pair

You can update the first node in an HA pair by initiating a takeover by the node’s partner. The partner serves the node’s data while the first node is upgraded.

If you are performing a major upgrade, the first node to be upgraded must be the same node on which you configured the data LIFs for external connectivity and installed the first ONTAP image.

After upgrading the first node, you should upgrade the partner node as quickly as possible. Do not allow the two nodes to remain in a state of version mismatch longer than necessary.

1. Update the first node in the cluster by invoking an AutoSupport message: `autosupport invoke -node * -type all -message "Starting_NDU"`

This AutoSupport notification includes a record of the system status just prior to update. It saves useful troubleshooting information in case there is a problem with the update process.

If the cluster is not configured to send AutoSupport messages, a copy of the notification is saved locally.

2. Set the privilege level to advanced, entering **y** when prompted to continue: `set -privilege advanced`

The advanced prompt (`*>`) appears.

3. Set the new ONTAP software image to be the default image: `system image modify {-node nodenameA -iscurrent false} -isdefault true`

The `system image modify` command uses an extended query to change the new ONTAP software image (which is installed as the alternate image) to the default image for the node.

4. Monitor the progress of the update: `system node upgrade-revert show`
5. Verify that the new ONTAP software image is set as the default image: `system image show`

In the following example, image2 is the new ONTAP version and is set as the default image on node0:

```
cluster1::*> system image show
Node      Image      Is      Is      Install
          Image    Default Current Version  Date
-----
node0
  image1  false    true    X.X.X   MM/DD/YYYY TIME
  image2  true     false   Y.Y.Y   MM/DD/YYYY TIME
node1
  image1  true     true    X.X.X   MM/DD/YYYY TIME
  image2  false   false   Y.Y.Y   MM/DD/YYYY TIME
4 entries were displayed.
```

6. Disable automatic giveback on the partner node if it is enabled: `storage failover modify -node nodenameB -auto-giveback false`

If the cluster is a two-node cluster, a message is displayed warning you that disabling automatic giveback prevents the management cluster services from going online in the event of an alternating-failure scenario. Enter `y` to continue.

7. Verify that automatic giveback is disabled for node's partner: `storage failover show -node nodenameB -fields auto-giveback`

```
cluster1::> storage failover show -node node1 -fields auto-giveback
node      auto-giveback
-----
node1     false
1 entry was displayed.
```

8. Run the following command twice to determine whether the node to be updated is currently serving any clients `system node run -node nodenameA -command uptime`

The `uptime` command displays the total number of operations that the node has performed for NFS, CIFS, FC, and iSCSI clients since the node was last booted. For each protocol, you must run the command twice to determine whether the operation counts are increasing. If they are increasing, the node is currently serving clients for that protocol. If they are not increasing, the node is not currently serving clients for that protocol.

NOTE: You should make a note of each protocol that has increasing client operations so that after the node is updated, you can verify that client traffic has resumed.

The following example shows a node with NFS, CIFS, FC, and iSCSI operations. However, the node is currently serving only NFS and iSCSI clients.

```

cluster1::> system node run -node node0 -command uptime
  2:58pm up 7 days, 19:16 800000260 NFS ops, 1017333 CIFS ops, 0 HTTP
ops, 40395 FCP ops, 32810 iSCSI ops

cluster1::> system node run -node node0 -command uptime
  2:58pm up 7 days, 19:17 800001573 NFS ops, 1017333 CIFS ops, 0 HTTP
ops, 40395 FCP ops, 32815 iSCSI ops

```

9. Migrate all of the data LIFs away from the node: `network interface migrate-all -node nodenameA`
10. Verify any LIFs that you migrated: `network interface show`

For more information about parameters you can use to verify LIF status, see the `network interface show man` page.

The following example shows that node0's data LIFs migrated successfully. For each LIF, the fields included in this example enable you to verify the LIF's home node and port, the current node and port to which the LIF migrated, and the LIF's operational and administrative status.

```

cluster1::> network interface show -data-protocol nfs|cifs -role data
-home-node node0 -fields home-node,curr-node,curr-port,home-port,status-
admin,status-oper
vserver lif      home-node home-port curr-node curr-port status-oper
status-admin
-----
-----
vs0      data001 node0      e0a      node1    e0a      up        up
vs0      data002 node0      e0b      node1    e0b      up        up
vs0      data003 node0      e0b      node1    e0b      up        up
vs0      data004 node0      e0a      node1    e0a      up        up
4 entries were displayed.

```

11. Initiate a takeover: `storage failover takeover -ofnode nodenameA`

Do not specify the `-option immediate` parameter, because a normal takeover is required for the node that is being taken over to boot onto the new software image. If you did not manually migrate the LIFs away from the node, they automatically migrate to the node's HA partner to ensure that there are no service disruptions.

The first node boots up to the Waiting for giveback state.

NOTE: If AutoSupport is enabled, an AutoSupport message is sent indicating that the node is out of cluster quorum. You can ignore this notification and proceed with the update.

12. Verify that the takeover is successful: `storage failover show`

You might see error messages indicating version mismatch and mailbox format problems. This is expected

behavior and it represents a temporary state in a major nondisruptive upgrade and is not harmful.

The following example shows that the takeover was successful. Node node0 is in the Waiting for giveback state, and its partner is in the In takeover state.

```
cluster1::> storage failover show
Node           Partner           Takeover
Possible State Description
-----
node0          node1              -          Waiting for giveback (HA
mailboxes)
node1          node0              false     In takeover
2 entries were displayed.
```

13. Wait at least eight minutes for the following conditions to take effect:

- Client multipathing (if deployed) is stabilized.
- Clients are recovered from the pause in an I/O operation that occurs during takeover.

The recovery time is client specific and might take longer than eight minutes, depending on the characteristics of the client applications.

14. Return the aggregates to the first node: `storage failover giveback -ofnode nodenameA`

The giveback first returns the root aggregate to the partner node and then, after that node has finished booting, returns the non-root aggregates and any LIFs that were set to automatically revert. The newly booted node begins to serve data to clients from each aggregate as soon as the aggregate is returned.

15. Verify that all aggregates have been returned: `storage failover show-giveback`

If the Giveback Status field indicates that there are no aggregates to give back, then all aggregates have been returned. If the giveback is vetoed, the command displays the giveback progress and which subsystem vetoed the giveback.

16. If any aggregates have not been returned, perform the following steps:

- a. Review the veto workaround to determine whether you want to address the “veto” condition or override the veto.

[High-availability configuration](#)

- b. If necessary, address the “veto” condition described in the error message, ensuring that any identified operations are terminated gracefully.
- c. Rerun the storage failover giveback command.

If you decided to override the “veto” condition, set the `-override-vetoes` parameter to true.

17. Wait at least eight minutes for the following conditions to take effect:

- Client multipathing (if deployed) is stabilized.
- Clients are recovered from the pause in an I/O operation that occurs during giveback.

The recovery time is client specific and might take longer than eight minutes, depending on the characteristics of the client applications.

18. Verify that the update was completed successfully for the node:

a. Go to the advanced privilege level: `set -privilege advanced`

b. Verify that update status is complete for the node: `system node upgrade-revert show -node nodenameA`

The status should be listed as complete.

If the status is not complete, from the node, run the `system node upgrade-revert upgrade` command. If the command does not complete the update, contact technical support.

c. Return to the admin privilege level: `set -privilege admin`

19. Verify that the node's ports are up: `network port show -node nodenameA`

You must run this command on a node that is upgraded to the higher version of ONTAP 9.

The following example shows that all of the node's ports are up:

```
cluster1::> network port show -node node0
```

(Mbps)						Speed
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper

node0						
	e0M	Default	-	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.

20. Revert the LIFs back to the node: `network interface revert *`

This command returns the LIFs that were migrated away from the node.

```
cluster1::> network interface revert *
8 entries were acted on.
```

21. Verify that the node's data LIFs successfully reverted back to the node, and that they are up: `network interface show`

The following example shows that all of the data LIFs hosted by the node have successfully reverted back to the node, and that their operational status is up:

```

cluster1::> network interface show
          Logical   Status   Network           Current
Current Is
Vserver  Interface  Admin/Oper  Address/Mask      Node      Port
Home
-----
vs0
          data001   up/up     192.0.2.120/24    node0     e0a
true
          data002   up/up     192.0.2.121/24    node0     e0b
true
          data003   up/up     192.0.2.122/24    node0     e0b
true
          data004   up/up     192.0.2.123/24    node0     e0a
true
4 entries were displayed.

```

22. If you previously determined that this node serves clients, verify that the node is providing service for each protocol that it was previously serving: `system node run -node nodenameA -command uptime`

The operation counts reset to zero during the update.

The following example shows that the updated node has resumed serving its NFS and iSCSI clients:

```

cluster1::> system node run -node node0 -command uptime
 3:15pm up 0 days, 0:16 129 NFS ops, 0 CIFS ops, 0 HTTP ops, 0 FCP
ops, 2 iSCSI ops

```

23. Reenable automatic giveback on the partner node if it was previously disabled: `storage failover modify -node nodenameB -auto-giveback true`

You should proceed to update the node's HA partner as quickly as possible. If you must suspend the update process for any reason, both nodes in the HA pair should be running the same ONTAP version.

Updating the partner node in an HA pair

After updating the first node in an HA pair, you update its partner by initiating a takeover on it. The first node serves the partner's data while the partner node is upgraded.

1. Set the privilege level to advanced, entering `y` when prompted to continue: `set -privilege advanced`

The advanced prompt (`*>`) appears.

2. Set the new ONTAP software image to be the default image: `system image modify {-node nodenameB -iscurrent false} -isdefault true`

The system image modify command uses an extended query to change the new ONTAP software image (which is installed as the alternate image) to be the default image for the node.

3. Monitor the progress of the update: `system node upgrade-revert show`
4. Verify that the new ONTAP software image is set as the default image: `system image show`

In the following example, `image2` is the new version of ONTAP and is set as the default image on the node:

```
cluster1::*> system image show
```

Node	Image	Is Default	Is Current	Version	Install Date

node0	image1	false	false	X.X.X	MM/DD/YYYY TIME
	image2	true	true	Y.Y.Y	MM/DD/YYYY TIME
node1	image1	false	true	X.X.X	MM/DD/YYYY TIME
	image2	true	false	Y.Y.Y	MM/DD/YYYY TIME

4 entries were displayed.

5. Disable automatic giveback on the partner node if it is enabled: `storage failover modify -node nodenameA -auto-giveback false`

If the cluster is a two-node cluster, a message is displayed warning you that disabling automatic giveback prevents the management cluster services from going online in the event of an alternating-failure scenario. Enter `y` to continue.

6. Verify that automatic giveback is disabled for the partner node: `storage failover show -node nodenameA -fields auto-giveback`

```
cluster1::> storage failover show -node node0 -fields auto-giveback
```

node	auto-giveback

node0	false

1 entry was displayed.

7. Run the following command twice to determine whether the node to be updated is currently serving any clients: `system node run -node nodenameB -command uptime`

The `uptime` command displays the total number of operations that the node has performed for NFS, CIFS, FC, and iSCSI clients since the node was last booted. For each protocol, you must run the command twice to determine whether the operation counts are increasing. If they are increasing, the node is currently serving clients for that protocol. If they are not increasing, the node is not currently serving clients for that protocol.

NOTE: You should make a note of each protocol that has increasing client operations so that after the node

is updated, you can verify that client traffic has resumed.

The following example shows a node with NFS, CIFS, FC, and iSCSI operations. However, the node is currently serving only NFS and iSCSI clients.

```
cluster1::> system node run -node node1 -command uptime
  2:58pm up 7 days, 19:16 800000260 NFS ops, 1017333 CIFS ops, 0 HTTP
ops, 40395 FCP ops, 32810 iSCSI ops

cluster1::> system node run -node node1 -command uptime
  2:58pm up 7 days, 19:17 800001573 NFS ops, 1017333 CIFS ops, 0 HTTP
ops, 40395 FCP ops, 32815 iSCSI ops
```

8. Migrate all of the data LIFs away from the node: `network interface migrate-all -node nodenameB`
9. Verify the status of any LIFs that you migrated: `network interface show`

For more information about parameters you can use to verify LIF status, see the `network interface show man` page.

The following example shows that node1's data LIFs migrated successfully. For each LIF, the fields included in this example enable you to verify the LIF's home node and port, the current node and port to which the LIF migrated, and the LIF's operational and administrative status.

```
cluster1::> network interface show -data-protocol nfs|cifs -role data
-home-node node1 -fields home-node,curr-node,curr-port,home-port,status-
admin,status-oper
vserver lif      home-node home-port curr-node curr-port status-oper
status-admin
-----
vs0      data001 node1      e0a      node0    e0a      up       up
vs0      data002 node1      e0b      node0    e0b      up       up
vs0      data003 node1      e0b      node0    e0b      up       up
vs0      data004 node1      e0a      node0    e0a      up       up
4 entries were displayed.
```

10. Initiate a takeover: `storage failover takeover -ofnode nodenameB -option allow-version-mismatch`

Do not specify the `-option immediate` parameter, because a normal takeover is required for the node that is being taken over to boot onto the new software image. If you did not manually migrate the LIFs away from the node, they automatically migrate to the node's HA partner so that there are no service disruptions.

The node that is taken over boots up to the `Waiting for giveback` state.

NOTE: If AutoSupport is enabled, an AutoSupport message is sent indicating that the node is out of cluster

quorum. You can ignore this notification and proceed with the update.

11. Verify that the takeover was successful: `storage failover show`

The following example shows that the takeover was successful. Node `node1` is in the `Waiting for giveback` state, and its partner is in the `In takeover` state.

```
cluster1::> storage failover show
Node           Partner           Takeover
-----
-----
node0          node1              -           In takeover
node1          node0              false      Waiting for giveback (HA
mailboxes)
2 entries were displayed.
```

12. Wait at least eight minutes for the following conditions to take effect:

- Client multipathing (if deployed) is stabilized.
- Clients are recovered from the pause in I/O that occurs during takeover.

The recovery time is client-specific and might take longer than eight minutes, depending on the characteristics of the client applications.

13. Return the aggregates to the partner node: `storage failover giveback -ofnode nodenameB`

The giveback operation first returns the root aggregate to the partner node and then, after that node has finished booting, returns the non-root aggregates and any LIFs that were set to automatically revert. The newly booted node begins to serve data to clients from each aggregate as soon as the aggregate is returned.

14. Verify that all aggregates are returned: `storage failover show-giveback`

If the `Giveback Status` field indicates that there are no aggregates to give back, then all aggregates are returned. If the giveback is vetoed, the command displays the giveback progress and which subsystem vetoed the giveback operation.

15. If any aggregates are not returned, perform the following steps:

- a. Review the veto workaround to determine whether you want to address the “veto” condition or override the veto.

[High-availability configuration](#)

- b. If necessary, address the “veto” condition described in the error message, ensuring that any identified operations are terminated gracefully.
- c. Rerun the `storage failover giveback` command.

If you decided to override the “veto” condition, set the `-override-vetoes` parameter to `true`.

16. Wait at least eight minutes for the following conditions to take effect:

- Client multipathing (if deployed) is stabilized.
- Clients are recovered from the pause in an I/O operation that occurs during giveback.

The recovery time is client specific and might take longer than eight minutes, depending on the characteristics of the client applications.

17. Verify that the update was completed successfully for the node:

- Go to the advanced privilege level: `set -privilege advanced`
- Verify that update status is complete for the node: `system node upgrade-revert show -node nodenameB`

The status should be listed as complete.

If the status is not complete, from the node, run the `system node upgrade-revert upgrade` command. If the command does not complete the update, contact technical support.

- Return to the admin privilege level: `set -privilege admin`

18. Verify that the node's ports are up: `network port show -node nodenameB`

You must run this command on a node that has been upgraded to ONTAP 9.4.

The following example shows that all of the node's data ports are up:

```
cluster1::> network port show -node node1
```

						Speed
(Mbps)	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper

node1						
	e0M	Default	-	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.

19. Revert the LIFs back to the node: `network interface revert *`

This command returns the LIFs that were migrated away from the node.

```
cluster1::> network interface revert *
8 entries were acted on.
```

20. Verify that the node's data LIFs successfully reverted back to the node, and that they are up: `network interface show`

The following example shows that all of the data LIFs hosted by the node is successfully reverted back to the node, and that their operational status is up:

```
cluster1::> network interface show
Current Is
Vserver      Logical      Status      Network      Current
Home         Interface   Admin/Oper  Address/Mask Node          Port
-----
-----
vs0
              data001      up/up       192.0.2.120/24 node1         e0a
true
              data002      up/up       192.0.2.121/24 node1         e0b
true
              data003      up/up       192.0.2.122/24 node1         e0b
true
              data004      up/up       192.0.2.123/24 node1         e0a
true
4 entries were displayed.
```

21. If you previously determined that this node serves clients, verify that the node is providing service for each protocol that it was previously serving: `system node run -node nodenameB -command uptime`

The operation counts reset to zero during the update.

The following example shows that the updated node has resumed serving its NFS and iSCSI clients:

```
cluster1::> system node run -node node1 -command uptime
3:15pm up 0 days, 0:16 129 NFS ops, 0 CIFS ops, 0 HTTP ops, 0 FCP
ops, 2 iSCSI ops
```

22. If this was the last node in the cluster to be updated, trigger an AutoSupport notification: `autosupport invoke -node * -type all -message "Finishing_NDU"`

This AutoSupport notification includes a record of the system status just prior to update. It saves useful troubleshooting information in case there is a problem with the update process.

If the cluster is not configured to send AutoSupport messages, a copy of the notification is saved locally.

23. Confirm that the new ONTAP software is running on both nodes of the HA pair: `system node image show`

In the following example, image2 is the updated version of ONTAP and is the default version on both nodes:

```

cluster1::*> system node image show
      Is      Is      Install
Node  Image  Default Current Version  Date
-----
node0
      image1 false  false  X.X.X  MM/DD/YYYY TIME
      image2 true   true   Y.Y.Y  MM/DD/YYYY TIME
node1
      image1 false  false  X.X.X  MM/DD/YYYY TIME
      image2 true   true   Y.Y.Y  MM/DD/YYYY TIME
4 entries were displayed.

```

24. Reenable automatic giveback on the partner node if it was previously disabled: `storage failover modify -node nodenameA -auto-giveback true`
25. Verify that the cluster is in quorum and that services are running by using the `cluster show` and `cluster ring show` (advanced privilege level) commands.

You must perform this step before upgrading any additional HA pairs.

26. Return to the admin privilege level: `set -privilege admin`

Upgrade any additional HA pairs.

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.