



ONTAP upgrade methods

ONTAP 9

NetApp
April 24, 2024

This PDF was generated from https://docs.netapp.com/us-en/ontap/upgrade/concept_upgrade_methods.html on April 24, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- ONTAP upgrade methods 1
 - ONTAP software upgrade methods 1
 - Automated nondisruptive ONTAP upgrade 4
 - Manual upgrades 13

ONTAP upgrade methods

ONTAP software upgrade methods

You can perform an automated upgrade of your ONTAP software using System Manager. Alternately, you can perform an automated or manual upgrade using the ONTAP command line interface (CLI). The method you use to upgrade ONTAP depends upon your configuration, your current ONTAP version, and the number of nodes in your cluster. NetApp recommends using System Manager to perform automated upgrades unless your configuration requires a different approach. For example, if you have a MetroCluster configuration with 4 nodes running ONTAP 9.3 or later, you should use System Manager to perform an automated upgrade (sometimes referred to as automated nondisruptive upgrade or ANDU). If you have a MetroCluster configuration with 8 nodes running ONTAP 9.2 or earlier, you should use the CLI to perform a manual upgrade.

An upgrade can be executed using the rolling upgrade process or the batch upgrade process. Both are nondisruptive.

For automated upgrades, ONTAP automatically installs the target ONTAP image on each node, validates the cluster components to ensure that the cluster can be upgraded nondisruptively, and then executes a batch or rolling upgrade in the background based on the number of nodes. For manual upgrades, the administrator manually confirms that each node in the cluster is ready for upgrade, then performs steps to execute a rolling upgrade.


ONTAP rolling upgrades

The rolling upgrade process is the default for clusters with fewer than 8 nodes. In the rolling upgrade process, a node is taken offline and upgraded while its partner takes over its storage. When the node upgrade is complete, the partner node gives control back to the original owning node, and the process is repeated on the partner node. Each additional HA pair is upgraded in sequence until all HA pairs are running the target release.

ONTAP batch upgrades

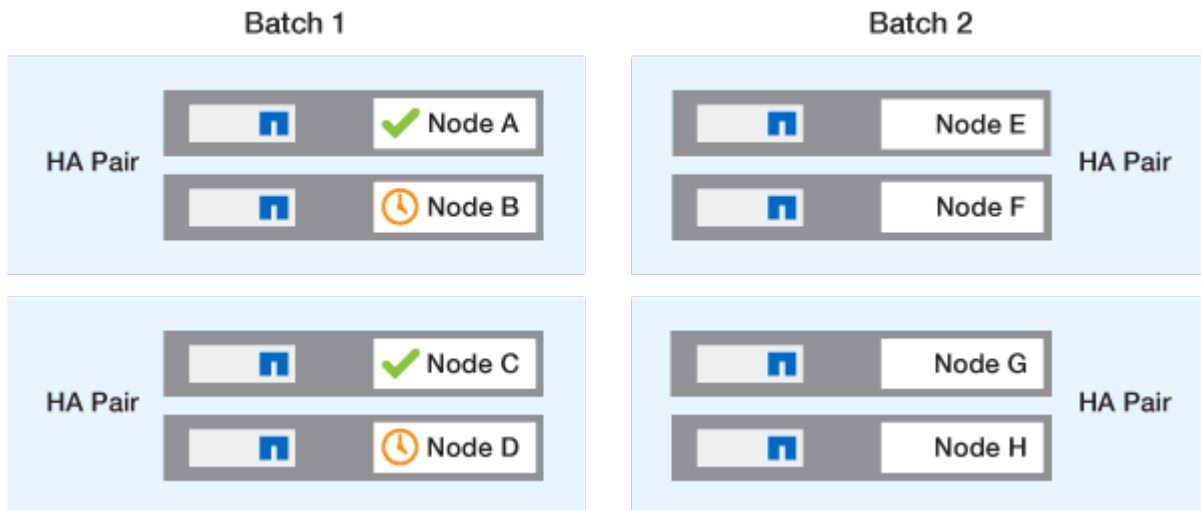
The batch upgrade process is the default for clusters of 8 nodes or more. In the batch upgrade process, the cluster is divided into two batches. Each batch contains multiple HA pairs. In the first batch, the first node of each HA pair is simultaneously upgraded with the first node of all other HA pairs in the batch.

In following example, there are two HA pairs in each batch. When the batch upgrade begins, Node A and Node C are upgraded simultaneously.



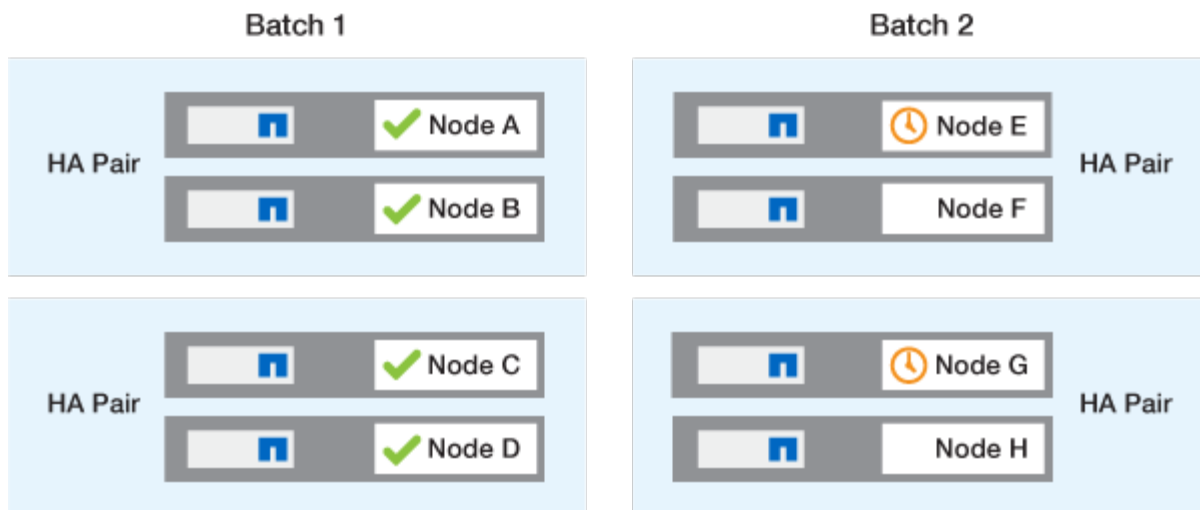
After the upgrade of the first nodes of each HA pair is complete, then the partner nodes in batch 1 are simultaneously upgraded.

In the following example, after Node A and Node C are upgraded, then Node B and Node D are simultaneously upgraded.




The process is then repeated for the nodes in batch 2; the first node of each HA pair is simultaneously upgraded with the first node of all other HA pairs in the batch.

In the following example, Node E and Node G are upgraded simultaneously.



After the upgrade of the first nodes of each HA pair is complete, then the partner nodes in batch 2 are simultaneously upgraded.

In the following example, Node F and Node H are simultaneously upgraded to complete the batch upgrade process.



Recommended ONTAP upgrade methods based on configuration

Upgrade methods supported by your configuration are listed in order of recommended usage.

| Configuration | ONTAP version | Number of nodes | Recommended upgrade method |
|---------------|---------------|-----------------|---|
| Standard | 9.0 or later | 2 or more | <ul style="list-style-type: none"> Automated nondisruptive using System Manager Automated nondisruptive using the CLI |
| Standard | 9.0 or later | Single | Automated disruptive |

| Configuration | ONTAP version | Number of nodes | Recommended upgrade method |
|---------------|----------------|-----------------|--|
| MetroCluster | 9.3 or later | 8 | <ul style="list-style-type: none"> Automated nondisruptive using the CLI Manual nondisruptive for 4 or 8 node MetroCluster using the CLI |
| MetroCluster | 9.3 or later | 2,4 | <ul style="list-style-type: none"> Automated nondisruptive using System Manager Automated nondisruptive using the CLI |
| MetroCluster | 9.2 or earlier | 4, 8 | Manual nondisruptive for 4 or 8 node MetroCluster using the CLI |
| MetroCluster | 9.2 or earlier | 2 | Manual nondisruptive for 2-node MetroCluster using the CLI |

ANDU using System Manager is the recommended upgrade method for all patch upgrades regardless of configuration.



A [manual disruptive upgrade](#) can be performed on any configuration. However, you should not perform a disruptive upgrade unless you can take the cluster offline for the duration of the upgrade. If you are operating in a SAN environment, you should be prepared to shut down or suspend all SAN clients before performing a disruptive upgrade. Disruptive upgrades are performed using the ONTAP CLI.

Automated nondisruptive ONTAP upgrade

When you perform an automated upgrade, ONTAP automatically installs the target ONTAP image on each node, validates that the cluster can be upgraded successfully, and then executes either a [batch or rolling upgrade](#) in the background based on the number of nodes in the cluster.

If it is supported by your configuration, you should use System Manager to perform an automated upgrade. If your configuration does not support automated upgrade using System Manager, you can use the ONTAP command line interface (CLI) to perform an automated upgrade.



Modifying the setting of the `storage failover modify-auto-giveback` command option before the start of an automatic nondisruptive upgrade (ANDU) has no impact on the upgrade process. The ANDU process ignores any preset value to this option during the takeover/giveback required for the update. For example, setting `-autogiveback` to `false` prior to beginning ANDU does not interrupt the automatic upgrade before giveback.

Before you begin

- You should [prepare for your upgrade](#).
- You should [download the ONTAP software image](#) for your target ONTAP release.

If you are performing a [direct multi-hop upgrade](#), you need to download both of the ONTAP images required for your specific [upgrade path](#).

- For each HA pair, each node should have one or more ports on the same broadcast domain.

If you have 8 or more nodes, the batch upgrade method is used in the automatic nondisruptive upgrade. In ONTAP 9.7 and earlier, if the batch method is used, LIFs are migrated to the HA partner of the node being upgraded. If the partners don't have any ports in the same broadcast domain, then the LIF migration fails.

In ONTAP 9.8 and later, if the batch method is used, LIFs are migrated to the other batch group.

- If you are upgrading ONTAP in a MetroCluster FC configuration, the cluster should be enabled for automatic unplanned switchover.
- If you don't plan to monitor the progress of the upgrade process, you should [request EMS notifications of errors that might require manual intervention](#).
- If you have a single-node cluster follow the [automated-disruptive upgrade](#) process.

Upgrades of single-node clusters are disruptive.

Example 1. Steps

System Manager


1. Validate the ONTAP target image:



If you are upgrading a MetroCluster configuration, you should validate Cluster A and then repeat the validation process on Cluster B.

- a. Depending on the ONTAP version that you are running, perform one of the following steps:

| If you are running... | Do this... |
|-------------------------|---|
| ONTAP 9.8 or later | Click Cluster > Overview . |
| ONTAP 9.5, 9.6, and 9.7 | Click Configuration > Cluster > Update . |
| ONTAP 9.4 or earlier | Click Configuration > Cluster Update . |

- b. In the right corner of the **Overview** pane, click .
- c. Click **ONTAP Update**.
- d. In the **Cluster Update** tab, add a new image or select an available image.

| If you want to... | Then... |
|---|---|
| Add a new software image from a local folder You should have already downloaded the image to the local client. | <ol style="list-style-type: none">i. Under Available Software Images, click Add from Local.ii. Browse to the location you saved the software image, select the image, and then click Open. |
| Add a new software image from an HTTP or FTP server | <ol style="list-style-type: none">i. Click Add from Server.ii. In the Add a New Software Image dialog box, enter the URL of the HTTP or FTP server to which you downloaded the ONTAP software image from the NetApp Support Site. For anonymous FTP, you must specify the URL in the ftp://anonymous@ftpserver format.iii. Click Add. |
| Select an available image | Choose one of the listed images. |

- e. Click **Validate** to run the pre-upgrade validation checks.

If any errors or warnings are found during validation, they are displayed along with a list of corrective actions. You must resolve all errors before proceeding with the upgrade. It is best

practice to also resolve warnings.

2. Click **Next**.
3. Click **Update**.

Validation is performed again. Any remaining errors or warnings are displayed along with a list of corrective actions. Errors must be corrected before you can proceed with the upgrade. If the validation is completed with warnings, you correct the warnings or choose **Update with warnings**.



By default, ONTAP uses the [batch upgrade process](#) to upgrade clusters with eight or more nodes. Beginning in ONTAP 9.10.1, if preferred, you can select **Update one HA pair at a time** to override the default and have your cluster upgrade one HA pair at a time using the rolling upgrade process.

For MetroCluster configurations with more than 2 nodes, the ONTAP upgrade process starts simultaneously on the HA pairs at both sites. For a 2-node MetroCluster configuration, the upgrade is started first on the site where the upgrade is not initiated. The upgrade on the remaining site begins after the first upgrade is fully completed.

4. If your upgrade pauses because of an error, click the error message to view the details, then correct the error and [resume the upgrade](#).

After you finish

After the upgrade is completed successfully, the node reboots, and you are redirected to the System Manager login page. If the node takes a long time to reboot, you should refresh your browser.

CLI

1. Validate the ONTAP target software image



If you are upgrading a MetroCluster configuration you should first execute the following steps on cluster A, then execute the same steps on cluster B.

- a. Delete the previous ONTAP software package:

```
cluster image package delete -version previous_ONTAP_Version
```

- b. Load the target ONTAP software image into the cluster package repository:

```
cluster image package get -url location
```

```
cluster1::> cluster image package get -url  
http://www.example.com/software/9.13.1/image.tgz
```

```
Package download completed.  
Package processing completed.
```

If you are performing a [direct multi-hop upgrade](#), you also need to load the software package for

the intermediate version of ONTAP required for your upgrade. For example, if you are upgrading from 9.8 to 9.13.1, you need to load the software package for ONTAP 9.12.1, and then use the same command to load the software package for 9.13.1.

- c. Verify that the software package is available in the cluster package repository:

```
cluster image package show-repository
```

```
cluster1::> cluster image package show-repository
Package Version  Package Build Time
-----
9.13.1          MM/DD/YYYY 10:32:15
```

- d. Execute the automated pre-upgrade checks:

```
cluster image validate -version package_version_number
```

If you are performing a [direct multi-hop upgrade](#), you only need to use the target ONTAP package for verification. You don't need to validate the intermediate upgrade image separately. For example, if you are upgrading from 9.8 to 9.13.1, use the 9.13.1 package for verification. You don't need to validate the 9.12.1 package separately.

```
cluster1::> cluster image validate -version 9.13.1

WARNING: There are additional manual upgrade validation checks
that must be performed after these automated validation checks
have completed...
```

- e. Monitor the progress of the validation:

```
cluster image show-update-progress
```

- f. Complete all required actions identified by the validation.
g. If you are upgrading a MetroCluster configuration, repeat the above steps on cluster B.

2. Generate a software upgrade estimate:

```
cluster image update -version package_version_number -estimate-only
```



If you are upgrading a MetroCluster configuration, you can run this command on either Cluster A or Cluster B. You don't need to run it on both clusters.

The software upgrade estimate displays details about each component to be updated, as well as the estimated duration of the upgrade.

3. Perform the software upgrade:

```
cluster image update -version package_version_number
```

- If you are performing a [direct multi-hop upgrade](#), use the target ONTAP version for the `package_version_number`. For example, if you are upgrading from ONTAP 9.8 to 9.13.1, use 9.13.1 as the `package_version_number`.
- By default, ONTAP uses the [batch upgrade process](#) to upgrade clusters with eight or more nodes. If preferred, you can use the `-force-rolling` parameter to override the default process and have your cluster upgraded one node at a time using the rolling upgrade process.
- After completing each takeover and giveback, the upgrade waits for 8 minutes to enable client applications to recover from the pause in I/O that occurs during the takeover and giveback. If your environment requires more or less time for client stabilization, you can use the `-stabilize -minutes` parameter to specify a different amount of stabilization time.
- For MetroCluster configurations with 4 nodes more, the automated upgrade starts simultaneously on the HA pairs at both sites. For a 2-node MetroCluster configuration, the upgrade starts on the site where the upgrade is not initiated. The upgrade on the remaining site begins after the first upgrade is fully completed.

```
cluster1::> cluster image update -version 9.13.1

Starting validation for this update. Please wait..

It can take several minutes to complete validation...

WARNING: There are additional manual upgrade validation checks...

Pre-update Check      Status      Error-Action
-----
...
20 entries were displayed

Would you like to proceed with update ? {y|n}: y
Starting update...

cluster-1::>
```

4. Display the cluster update progress:

```
cluster image show-update-progress
```

If you are upgrading a 4-node or 8-node MetroCluster configuration, the `cluster image show-update-progress` command only displays the progress for the node on which you run the command. You must run the command on each node to see individual node progress.

5. Verify that the upgrade was completed successfully on each node.

```
cluster image show-update-progress
```

```
cluster1::> cluster image show-update-progress
```

| Elapsed | | Estimated |
|--------------------|-----------|-----------|
| Update Phase | Status | Duration |
| Duration | | |
| ----- | ----- | ----- |
| Pre-update checks | completed | 00:10:00 |
| 00:02:07 | | |
| Data ONTAP updates | completed | 01:31:00 |
| 01:39:00 | | |
| Post-update checks | completed | 00:10:00 |
| 00:02:00 | | |

3 entries were displayed.

Updated nodes: node0, node1.

6. Trigger an AutoSupport notification:

```
autosupport invoke -node * -type all -message "Finishing_NDU"
```

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

7. If you are upgrading a 2-node MetroCluster FC configuration, verify that the cluster is enabled for automatic unplanned switchover.



If you are upgrading a standard configuration, a MetroCluster IP configuration, or a MetroCluster FC configuration greater than 2 nodes, you don't need to perform this step.

- a. Check whether automatic unplanned switchover is enabled:

```
metrocluster show
```

If automatic unplanned switchover is enabled, the following statement appears in the command output:

```
AUSO Failure Domain      auso-on-cluster-disaster
```

- b. If the statement does not appear in the output, enable automatic unplanned switchover:

```
metrocluster modify -auto-switchover-failure-domain auso-on-  
cluster-disaster
```

- c. Verify that automatic unplanned switchover has been enabled:

```
metrocluster show
```

Resume ONTAP software upgrade after an error in the automated upgrade process

If an automated ONTAP software upgrade pauses because of an error, you should resolve the error and then continue the upgrade. After the error is resolved, you can choose to continue the automated upgrade process or complete the upgrade process manually. If you choose to continue the automated upgrade, don't perform any of the upgrade steps manually.

Example 2. Steps

System Manager

1. Depending on the ONTAP version that you are running, perform one of the following steps:

| If you are running... | Then... |
|------------------------|--|
| ONTAP 9.8 or later | Click Cluster > Overview |
| ONTAP 9.7, 9.6, or 9.5 | Click Configuration > Cluster > Update . |
| ONTAP 9.4 or earlier | <ul style="list-style-type: none">• Click Configuration > Cluster Update.• In the right corner of the Overview pane, click the three blue vertical dots, and select ONTAP Update. |

2. Continue the automated upgrade or cancel it and continue manually.

| If you want to... | Then... |
|--|-----------------------|
| Resume the automated upgrade | Click Resume . |
| Cancel the automated upgrade and continue manually | Click Cancel . |

CLI

1. View the upgrade error:

```
cluster image show-update-progress
```

2. Resolve the error.
3. Resume the upgrade:

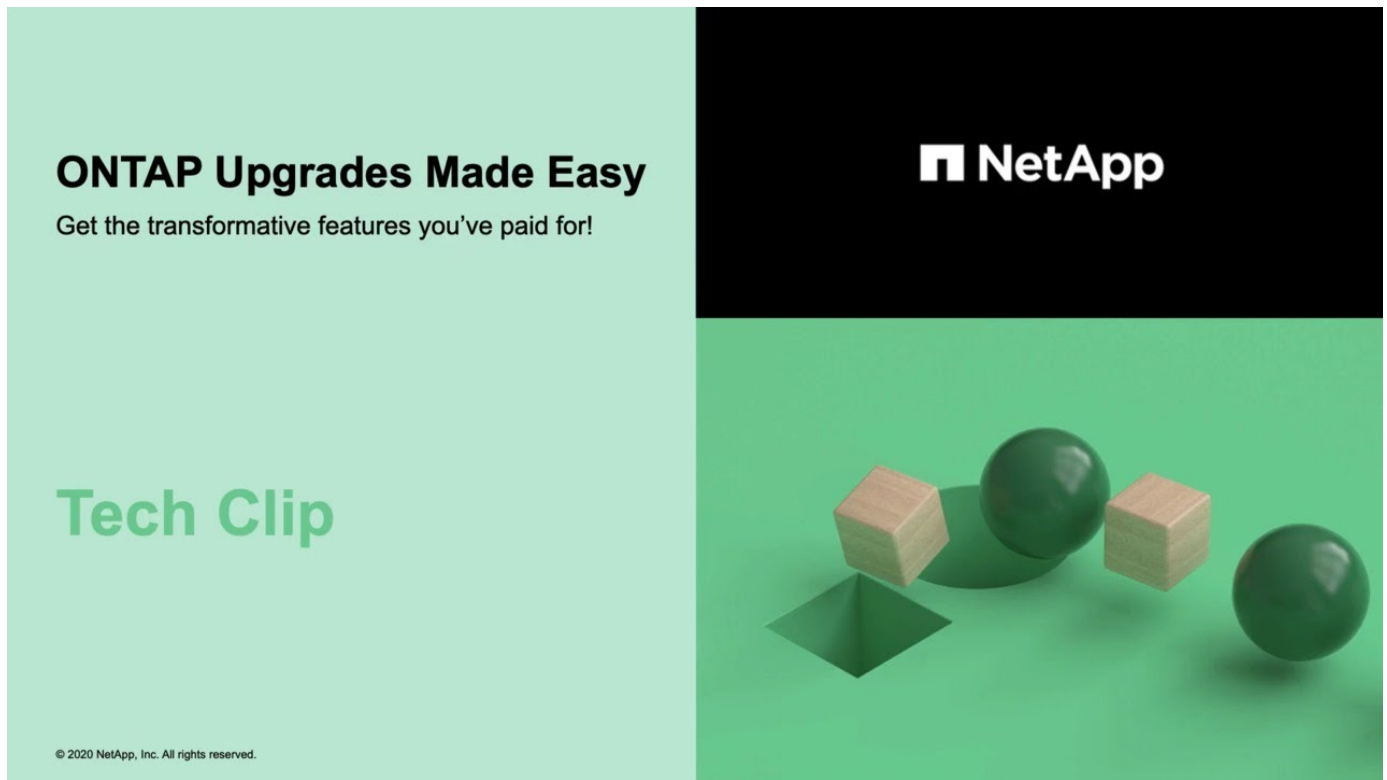
| If you want to... | Enter the following command... |
|--|--|
| Resume the automated upgrade | <pre>cluster image resume-update</pre> |
| Cancel the automated upgrade and continue manually | <pre>cluster image cancel-update</pre> |

After you finish

[Perform post-upgrade checks.](#)

Video: Upgrades made easy

Take a look at the simplified ONTAP upgrade capabilities of System Manager in ONTAP 9.8.



Related information

- [Launch Active IQ](#)
- [Active IQ documentation](#)

Manual upgrades

Install the ONTAP software package for manual upgrades

After you have downloaded the ONTAP software package for a manual upgrade, you must install it locally before you begin your upgrade.

Steps

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

The advanced prompt (***>**) appears.

2. Install the image.

| If you have the following configuration... | Use this command... |
|--|---|
| <ul style="list-style-type: none"> • Non-MetroCluster • 2-node MetroCluster | <pre data-bbox="873 191 1430 344">system node image update -node * -package _location_ -replace -package true -setdefault true -background true</pre> <p data-bbox="841 411 1481 512"><i>location</i> can be a web server or a local folder, depending on the ONTAP version. See the <code>system node image update</code> man page for details.</p> <p data-bbox="841 552 1471 688">This command installs the software image on all of the nodes simultaneously. To install the image on each node one at a time, do not specify the <code>-background</code> parameter.</p> |
| <ul style="list-style-type: none"> • 4-node MetroCluster • 8-node MetroCluster configuration | <pre data-bbox="873 772 1430 926">system node image update -node * -package location -replace -package true -background true -setdefault false</pre> <p data-bbox="841 993 1422 1020">You must issue this command on both clusters.</p> <p data-bbox="841 1060 1468 1161">This command uses an extended query to change the target software image, which is installed as the alternate image on each node.</p> |

3. Enter `y` to continue when prompted.
4. Verify that the software image is installed on each node.

```
system node image show-update-progress -node *
```

This command displays the current status of the software image installation. You should continue to run this command until all nodes report a **Run Status** of **Exited**, and an **Exit Status** of **Success**.

The `system node image update` command can fail and display error or warning messages. After resolving any errors or warnings, you can run the command again.

This example shows a two-node cluster in which the software image is installed successfully on both nodes:


```

cluster1::*> system node image show-update-progress -node *
There is no update/install in progress
Status of most recent operation:
    Run Status:      Exited
    Exit Status:     Success
    Phase:           Run Script
    Exit Message:    After a clean shutdown, image2 will be set as
the default boot image on node0.
There is no update/install in progress
Status of most recent operation:
    Run Status:      Exited
    Exit Status:     Success
    Phase:           Run Script
    Exit Message:    After a clean shutdown, image2 will be set as
the default boot image on node1.
2 entries were acted on.

```

Manual nondisruptive ONTAP upgrade using the CLI (standard configurations)

Automated upgrade using System Manager is the preferred upgrade method. If System Manager does not support your configuration, you can use the ONTAP command line interface (CLI) to perform a manual nondisruptive upgrade. To upgrade a cluster of two or more nodes using the manual nondisruptive method, you must initiate a failover operation on each node in an HA pair, update the “failed” node, initiate giveback, and then repeat the process for each HA pair in the cluster.

Before you begin

You must have satisfied upgrade [preparation](#) requirements.

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 [mixed version](#) state longer than necessary.

Steps

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 Default | Is Current | Version | Install 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, SMB, 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.



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, SMB, 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.



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.
- 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, 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
```

| | | | | | | Speed |
|---------------------------|-------|---------|------------------|-------|-------|------------|
| (Mbps) | | | | | | |
| 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, SMB, 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, SMB, 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.

A warning is displayed. You must enter `y` to continue.

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



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 Possible | State Description |
|-------|---------|----------------------|--|
| 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:

- Review the veto workaround to determine whether you want to address the “veto” condition or override the veto.
- If necessary, address the “veto” condition described in the error message, ensuring that any identified operations are terminated gracefully.
- 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:

- 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 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.

- c. 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) | | | | | | |
| Node | 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
```

| | Logical | Status | Network | Current | |
|------------|-----------|------------|----------------|---------|-------|
| Current Is | | | | | |
| Vserver | Interface | Admin/Oper | Address/Mask | Node | Port |
| Home | | | | | |
| ----- | ----- | ----- | ----- | ----- | ----- |
| 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:

```
set -privilege advanced
```

```
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
```

| 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 | 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
```

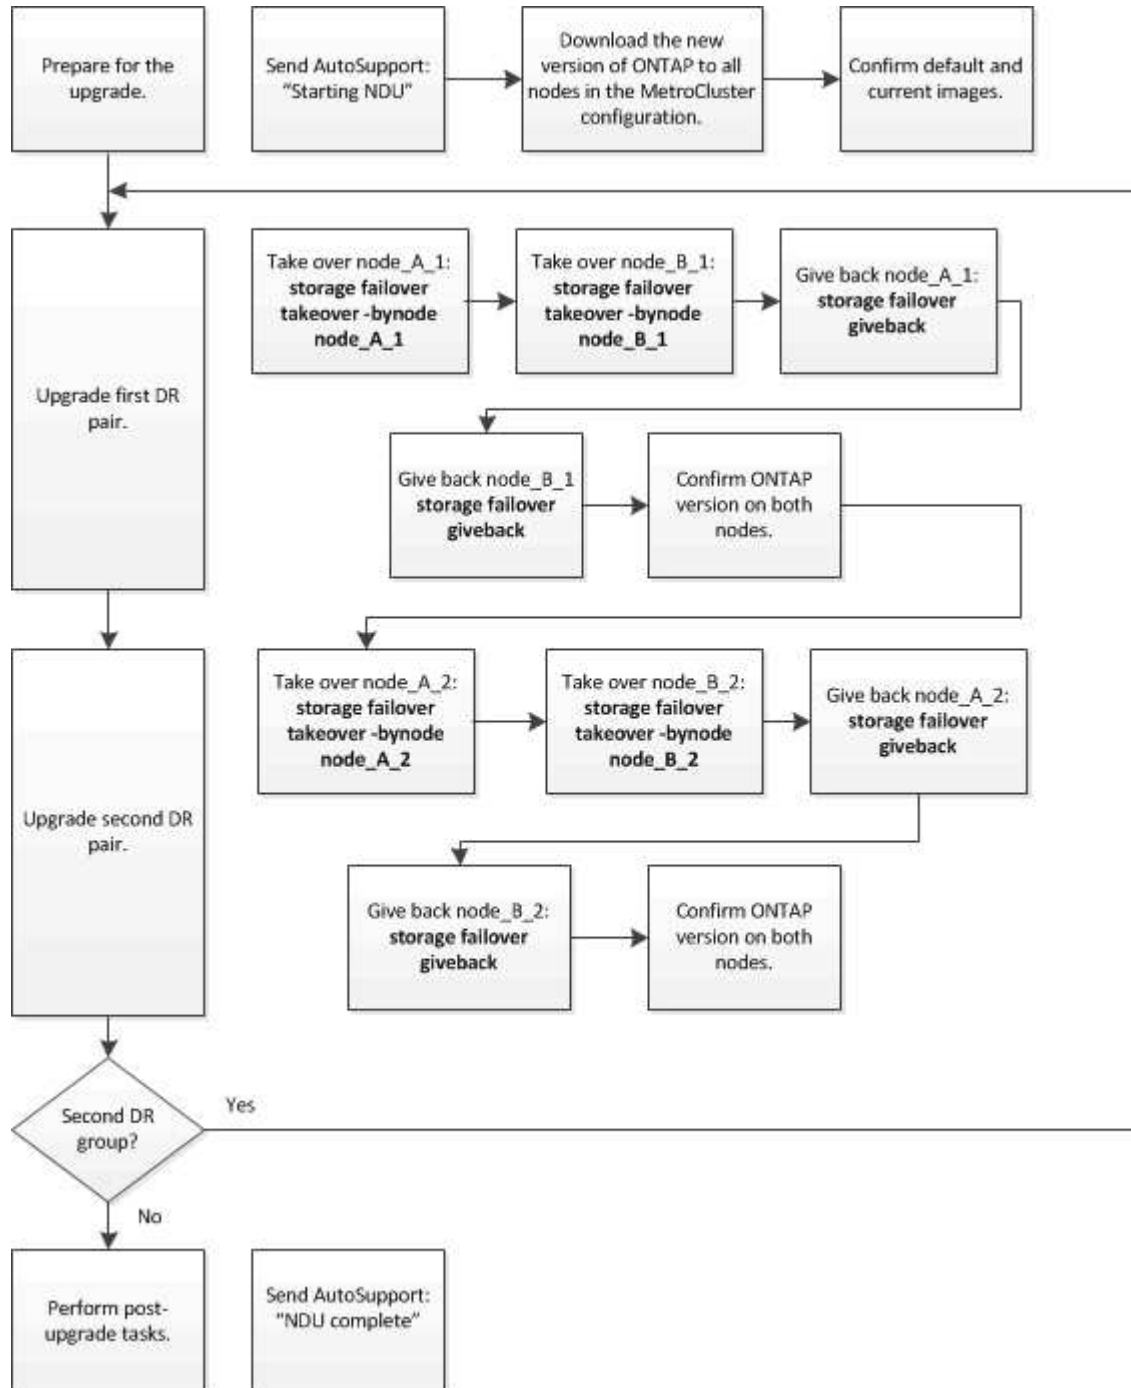
27. Upgrade any additional HA pairs.

Manual nondisruptive ONTAP upgrade of a four- or eight-node MetroCluster configuration using the CLI

A manual upgrade of a four- or eight-node MetroCluster configuration involves preparing for the update, updating the DR pairs in each of the one or two DR groups simultaneously, and performing post-upgrade tasks.

- This task applies to the following configurations:


- Four-node MetroCluster FC or IP configurations running ONTAP 9.2 or earlier
- Eight-node MetroCluster FC configurations, regardless of ONTAP version
- If you have a two-node MetroCluster configuration, do not use this procedure.
- The following tasks refer to the old and new versions of ONTAP.
 - When upgrading, the old version is a previous version of ONTAP, with a lower version number than the new version of ONTAP.
 - When downgrading, the old version is a later version of ONTAP, with a higher version number than the new version of ONTAP.
- This task uses the following high-level workflow:



Differences when updating ONTAP software on an eight-node or four-node MetroCluster configuration

The MetroCluster software upgrade process differs, depending on whether there are eight or four nodes in the MetroCluster configuration.

A MetroCluster configuration consists of one or two DR groups. Each DR group consists of two HA pairs, one HA pair at each MetroCluster cluster. An eight-node MetroCluster includes two DR groups:



You upgrade one DR group at a time.

For four-node MetroCluster configurations:

1. Upgrade DR Group One:
 - a. Upgrade node_A_1 and node_B_1.
 - b. Upgrade node_A_2 and node_B_2.

For eight-node MetroCluster configurations, you perform the DR group upgrade procedure twice:

1. Upgrade DR Group One:
 - a. Upgrade node_A_1 and node_B_1.
 - b. Upgrade node_A_2 and node_B_2.
2. Upgrade DR Group Two:
 - a. Upgrade node_A_3 and node_B_3.
 - b. Upgrade node_A_4 and node_B_4.


Preparing to upgrade a MetroCluster DR group

Before you upgrade the ONTAP software on the nodes, you must identify the DR relationships among the nodes, send an AutoSupport message that you are initiating an upgrade, and confirm the ONTAP version running on each node.

You must have [downloaded](#) and [installed](#) the software images.

This task must be repeated on each DR group. If the MetroCluster configuration consists of eight nodes, there are two DR groups. Thereby, this task must be repeated on each DR group.

The examples provided in this task use the names shown in the following illustration to identify the clusters and nodes:



1. Identify the DR pairs in the configuration:

```
metrocluster node show -fields dr-partner
```

```
cluster_A::> metrocluster node show -fields dr-partner
(metrocluster node show)
dr-group-id cluster      node      dr-partner
-----
1           cluster_A    node_A_1  node_B_1
1           cluster_A    node_A_2  node_B_2
1           cluster_B    node_B_1  node_A_1
1           cluster_B    node_B_2  node_A_2
4 entries were displayed.

cluster_A::>
```

2. Set the privilege level from admin to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (*>) appears.

3. Confirm the ONTAP version on cluster_A:

```
system image show
```

```
cluster_A::*> system image show
```

| Node | Image | Is Default | Is Current | Version | Install Date |
|----------|--------|------------|------------|---------|-----------------|
| node_A_1 | image1 | true | true | X.X.X | MM/DD/YYYY TIME |
| | image2 | false | false | Y.Y.Y | MM/DD/YYYY TIME |
| node_A_2 | image1 | true | true | X.X.X | MM/DD/YYYY TIME |
| | image2 | false | false | Y.Y.Y | MM/DD/YYYY TIME |

```
4 entries were displayed.

cluster_A::>
```

4. Confirm the version on cluster_B:

```
system image show
```

```
cluster_B::*> system image show
```

| Node | Image | Is Default | Is Current | Version | Install Date |
|----------|--------|---------------|---------------|---------|-----------------|
| ----- | | | | | |
| node_B_1 | | | | | |
| | image1 | true | true | X.X.X | MM/DD/YYYY TIME |
| | image2 | false | false | Y.Y.Y | MM/DD/YYYY TIME |
| node_B_2 | | | | | |
| | image1 | true | true | X.X.X | MM/DD/YYYY TIME |
| | image2 | false | false | Y.Y.Y | MM/DD/YYYY TIME |

4 entries were displayed.

```
cluster_B::>
```

5. Trigger an AutoSupport notification:

```
autosupport invoke -node * -type all -message "Starting_NDU"
```

This AutoSupport notification includes a record of the system status before the upgrade. It saves useful troubleshooting information if there is a problem with the upgrade process.

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

6. For each node in the first set, set the target ONTAP software image to be the default image:

```
system image modify {-node nodename -iscurrent false} -isdefault true
```

This command uses an extended query to change the target software image, which is installed as the alternate image, to be the default image for the node.

7. Verify that the target ONTAP software image is set as the default image on cluster_A:

```
system image show
```

In the following example, image2 is the new ONTAP version and is set as the default image on each of the nodes in the first set:

```
cluster_A::*> system image show
```

| Node | Image | Is Default | Is Current | Version | Install Date |
|----------|--------|---------------|---------------|---------|-----------------|
| ----- | | | | | |
| node_A_1 | image1 | false | true | X.X.X | MM/DD/YYYY TIME |
| | image2 | true | false | Y.Y.Y | MM/DD/YYYY TIME |
| node_A_2 | image1 | false | true | X.X.X | MM/DD/YYYY TIME |
| | image2 | true | false | Y.Y.Y | MM/DD/YYYY TIME |

2 entries were displayed.

- a. Verify that the target ONTAP software image is set as the default image on cluster_B:

```
system image show
```

The following example shows that the target version is set as the default image on each of the nodes in the first set:

```
cluster_B::*> system image show
```

| Node | Image | Is Default | Is Current | Version | Install Date |
|----------|--------|---------------|---------------|---------|-----------------|
| ----- | | | | | |
| node_A_1 | image1 | false | true | X.X.X | MM/DD/YYYY TIME |
| | image2 | true | false | Y.Y.Y | MM/YY/YYYY TIME |
| node_A_2 | image1 | false | true | X.X.X | MM/DD/YYYY TIME |
| | image2 | true | false | Y.Y.Y | MM/DD/YYYY TIME |

2 entries were displayed.

8. Determine whether the nodes to be upgraded are currently serving any clients twice for each node:

```
system node run -node target-node -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 need to 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.



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

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

```
cluster_x::> 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

cluster_x::> 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
```

Updating the first DR pair in a MetroCluster DR group

You must perform a takeover and giveback of the nodes in the correct order to make the new version of ONTAP the current version of the node.

All nodes must be running the old version of ONTAP.

In this task, node_A_1 and node_B_1 are upgraded.

If you have upgraded the ONTAP software on the first DR group, and are now upgrading the second DR group in an eight-node MetroCluster configuration, in this task you would be updating node_A_3 and node_B_3.

1. If MetroCluster Tiebreaker software is enabled, disabled it.
2. For each node in the HA pair, disable automatic giveback:

```
storage failover modify -node target-node -auto-giveback false
```

This command must be repeated for each node in the HA pair.

3. Verify that automatic giveback is disabled:

```
storage failover show -fields auto-giveback
```

This example shows that automatic giveback has been disabled on both nodes:

```
cluster_x::> storage failover show -fields auto-giveback
node      auto-giveback
-----
node_x_1  false
node_x_2  false
2 entries were displayed.
```

4. Ensure that I/O is not exceeding ~50% for each controller and that CPU utilization is not exceeding ~50% per controller.
5. Initiate a takeover of the target node on cluster_A:

Do not specify the `-option immediate` parameter, because a normal takeover is required for the nodes that are being taken over to boot onto the new software image.

- a. Take over the DR partner on cluster_A (node_A_1):

```
storage failover takeover -ofnode node_A_1
```

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



If AutoSupport is enabled, then an AutoSupport message is sent indicating that the nodes are out of cluster quorum. You can ignore this notification and proceed with the upgrade.

- b. Verify that the takeover is successful:

```
storage failover show
```

The following example shows that the takeover is successful. Node_A_1 is in the "Waiting for giveback" state and node_A_2 is in the "In takeover" state.

```
cluster1::> storage failover show
```

| Node | Partner | Takeover Possible | State Description |
|----------|----------|----------------------|-------------------------------------|
| node_A_1 | node_A_2 | - | Waiting for giveback (HA mailboxes) |
| node_A_2 | node_A_1 | false | In takeover |

2 entries were displayed.

6. Take over the DR partner on cluster_B (node_B_1):

Do not specify the `-option immediate` parameter, because a normal takeover is required for the nodes that

are being taken over to boot onto the new software image.

a. Take over node_B_1:

```
storage failover takeover -ofnode node_B_1
```

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



If AutoSupport is enabled, then an AutoSupport message is sent indicating that the nodes are out of cluster quorum. You can ignore this notification and proceed with the upgrade.

b. Verify that the takeover is successful:

```
storage failover show
```

The following example shows that the takeover is successful. Node_B_1 is in the "Waiting for giveback" state and node_B_2 is in the "In takeover" state.

```
cluster1::> storage failover show
```

| Node | Partner | Takeover Possible | State Description |
|----------|----------|----------------------|--|
| node_B_1 | node_B_2 | - | Waiting for giveback (HA mailboxes) |
| node_B_2 | node_B_1 | false | In takeover |

2 entries were displayed.

7. Wait at least eight minutes to ensure the following conditions:

- 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.

8. Return the aggregates to the target nodes:

After upgrading MetroCluster IP configurations to ONTAP 9.5 or later, the aggregates will be in a degraded state for a short period before resynchronizing and returning to a mirrored state.

a. Give back the aggregates to the DR partner on cluster_A:

```
storage failover giveback -ofnode node_A_1
```

- b. Give back the aggregates to the DR partner on cluster_B:

```
storage failover giveback -ofnode node_B_1
```

The giveback operation first returns the root aggregate to the node and then, after the node has finished booting, returns the non-root aggregates.

9. Verify that all aggregates have been returned by issuing the following command on both clusters:

```
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.

10. If any aggregates have not been returned, do the following:
- Review the veto workaround to determine whether you want to address the “veto” condition or override the veto.
 - If necessary, address the “veto” condition described in the error message, ensuring that any identified operations are terminated gracefully.
 - Reenter the storage failover giveback command.

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

11. Wait at least eight minutes to ensure the following conditions:

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

12. Set the privilege level from admin to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (`*>`) appears.

13. Confirm the version on cluster_A:

```
system image show
```

The following example shows that System image2 should be the default and current version on node_A_1:

```
cluster_A::*> system image show
```

| Node | Image | Is Default | Is Current | Version | Install Date |
|----------|--------|---------------|---------------|---------|-----------------|
| ----- | | | | | |
| node_A_1 | | | | | |
| | image1 | false | false | X.X.X | MM/DD/YYYY TIME |
| | image2 | true | true | Y.Y.Y | MM/DD/YYYY TIME |
| node_A_2 | | | | | |
| | image1 | false | true | X.X.X | MM/DD/YYYY TIME |
| | image2 | true | false | Y.Y.Y | MM/DD/YYYY TIME |

4 entries were displayed.

```
cluster_A::>
```

14. Confirm the version on cluster_B:

```
system image show
```

The following example shows that System image2 (ONTAP 9.0.0) is the default and current version on node_A_1:

```
cluster_A::*> system image show
```

| Node | Image | Is Default | Is Current | Version | Install Date |
|----------|--------|---------------|---------------|---------|-----------------|
| ----- | | | | | |
| node_B_1 | | | | | |
| | image1 | false | false | X.X.X | MM/DD/YYYY TIME |
| | image2 | true | true | Y.Y.Y | MM/DD/YYYY TIME |
| node_B_2 | | | | | |
| | image1 | false | true | X.X.X | MM/DD/YYYY TIME |
| | image2 | true | false | Y.Y.Y | MM/DD/YYYY TIME |

4 entries were displayed.

```
cluster_A::>
```

Updating the second DR pair in a MetroCluster DR group

You must perform a takeover and giveback of the node in the correct order to make the new version of ONTAP the current version of the node.

You should have upgraded the first DR pair (node_A_1 and node_B_1).

In this task, node_A_2 and node_B_2 are upgraded.

If you have upgraded the ONTAP software on the first DR group, and are now updating the second DR group

in an eight-node MetroCluster configuration, in this task you are updating node_A_4 and node_B_4.

1. Migrate all of the data LIFs away from the node:

```
network interface migrate-all -node nodenameA
```

2. Initiate a takeover of the target node on cluster_A:

Do not specify the `-option immediate` parameter, because a normal takeover is required for the nodes that are being taken over to boot onto the new software image.

- a. Take over the DR partner on cluster_A:

```
storage failover takeover -ofnode node_A_2 -option allow-version-  
mismatch
```



The `allow-version-mismatch` option is not required for upgrades from ONTAP 9.0 to ONTAP 9.1 or for any patch upgrades.

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

If AutoSupport is enabled, then an AutoSupport message is sent indicating that the nodes are out of cluster quorum. You can ignore this notification and proceed with the upgrade.

- b. Verify that the takeover is successful:

```
storage failover show
```

The following example shows that the takeover is successful. Node_A_2 is in the "Waiting for giveback" state and node_A_1 is in the "In takeover" state.

```
cluster1::> storage failover show
```

| Node | Partner | Takeover Possible | State Description |
|----------|----------|----------------------|--|
| node_A_1 | node_A_2 | false | In takeover |
| node_A_2 | node_A_1 | - | Waiting for giveback (HA mailboxes) |

2 entries were displayed.

3. Initiate a takeover of the target node on cluster_B:

Do not specify the `-option immediate` parameter, because a normal takeover is required for the nodes that are being taken over to boot onto the new software image.

a. Take over the DR partner on cluster_B (node_B_2):

| If you are upgrading from... | Enter this command... |
|-------------------------------|---|
| ONTAP 9.2 or ONTAP 9.1 | <pre>storage failover takeover -ofnode node_B_2</pre> |
| ONTAP 9.0 or Data ONTAP 8.3.x | <pre>storage failover takeover -ofnode node_B_2 -option allow- version-mismatch</pre> <div><p>The <code>allow-version-mismatch</code> option is not required for upgrades from ONTAP 9.0 to ONTAP 9.1 or for any patch upgrades.</p></div> |

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



If AutoSupport is enabled, an AutoSupport message is sent indicating that the nodes are out of cluster quorum. You can safely ignore this notification and proceed with the upgrade.

b. Verify that the takeover is successful:

```
storage failover show
```

The following example shows that the takeover is successful. Node_B_2 is in the "Waiting for giveback" state and node_B_1 is in the "In takeover" state.

```
cluster1::> storage failover show

Node           Partner           Takeover
Possible State Description
-----
node_B_1       node_B_2           false    In takeover
node_B_2       node_B_1           -        Waiting for giveback (HA
mailboxes)
2 entries were displayed.
```

4. Wait at least eight minutes to ensure the following conditions:

- 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.

5. Return the aggregates to the target nodes:

After upgrading MetroCluster IP configurations to ONTAP 9.5, the aggregates will be in a degraded state for a short period before resynchronizing and returning to a mirrored state.

a. Give back the aggregates to the DR partner on cluster_A:

```
storage failover giveback -ofnode node_A_2
```

b. Give back the aggregates to the DR partner on cluster_B:

```
storage failover giveback -ofnode node_B_2
```

The giveback operation first returns the root aggregate to the node and then, after the node has finished booting, returns the non-root aggregates.

6. Verify that all aggregates have been returned by issuing the following command on both clusters:

```
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.

7. If any aggregates have not been returned, do the following:

- a. Review the veto workaround to determine whether you want to address the “veto” condition or override the veto.
- b. If necessary, address the “veto” condition described in the error message, ensuring that any identified operations are terminated gracefully.
- c. Reenter the storage failover giveback command.

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

8. Wait at least eight minutes to ensure the following conditions:

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

9. Set the privilege level from admin to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (*>) appears.

10. Confirm the version on cluster_A:

```
system image show
```

The following example shows that System image2 (target ONTAP image) is the default and current version on node_A_2:

```
cluster_B::*> system image show
```

| Node | Image | Is Default | Is Current | Version | Install Date |
|----------|--------|---------------|---------------|---------|-----------------|
| node_A_1 | image1 | false | false | X.X.X | MM/DD/YYYY TIME |
| | image2 | true | true | Y.Y.Y | MM/DD/YYYY TIME |
| node_A_2 | image1 | false | false | X.X.X | MM/DD/YYYY TIME |
| | image2 | true | true | Y.Y.Y | MM/DD/YYYY TIME |

4 entries were displayed.

```
cluster_A::>
```

11. Confirm the version on cluster_B:

```
system image show
```

The following example shows that System image2 (target ONTAP image) is the default and current version on node_B_2:

```
cluster_B::*> system image show
```

| Node | Image | Is Default | Is Current | Version | Install Date |
|----------|--------|---------------|---------------|---------|-----------------|
| node_B_1 | image1 | false | false | X.X.X | MM/DD/YYYY TIME |
| | image2 | true | true | Y.Y.Y | MM/DD/YYYY TIME |
| node_B_2 | image1 | false | false | X.X.X | MM/DD/YYYY TIME |
| | image2 | true | true | Y.Y.Y | MM/DD/YYYY TIME |

4 entries were displayed.

```
cluster_A::>
```

12. For each node in the HA pair, enable automatic giveback:

```
storage failover modify -node target-node -auto-giveback true
```

This command must be repeated for each node in the HA pair.

13. Verify that automatic giveback is enabled:

```
storage failover show -fields auto-giveback
```

This example shows that automatic giveback has been enabled on both nodes:

```
cluster_x::> storage failover show -fields auto-giveback
node      auto-giveback
-----
node_x_1  true
node_x_2  true
2 entries were displayed.
```

Nondisruptive upgrade of a two-node MetroCluster configuration in ONTAP 9.2 or earlier

How you upgrade a two-node MetroCluster configuration varies based on your ONTAP version. If you are running ONTAP 9.2 or earlier, you should use this procedure to perform a manual nondisruptive upgrade which includes initiating a negotiated switchover, updating the cluster at the “failed” site, initiating switchback, and then repeating the process on the cluster at the other site.

If you have a two-node MetroCluster configuration running ONTAP 9.3 or later, perform an [automated upgrade using System Manager](#).

Steps

1. Set the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (***>**) appears.

2. On the cluster to be upgraded, install the new ONTAP software image as the default:

```
system node image update -package package_location -setdefault true
-replace-package true
```



```
cluster_B::*> system node image update -package
http://www.example.com/NewImage.tgz -setdefault true -replace-package
true
```

3. Verify that the target software image is set as the default image:

```
system node image show
```

The following example shows that NewImage is set as the default image:

```
cluster_B::*> system node image show
```

| Node | Image | Is Default | Is Current | Version | Install Date |
|----------|----------|---------------|---------------|---------|-----------------|
| node_B_1 | OldImage | false | true | X.X.X | MM/DD/YYYY TIME |
| | NewImage | true | false | Y.Y.Y | MM/DD/YYYY TIME |

2 entries were displayed.

4. If the target software image is not set as the default image, then change it:

```
system image modify {-node * -iscurrent false} -isdefault true
```

5. Verify that all cluster SVMs are in a health state:

```
metrocluster vserver show
```

6. On the cluster that is not being updated, initiate a negotiated switchover:

```
metrocluster switchover
```

The operation can take several minutes. You can use the metrocluster operation show command to verify that the switchover is completed.

In the following example, a negotiated switchover is performed on the remote cluster ("cluster_A"). This causes the local cluster ("cluster_B") to halt so that you can update it.

```
cluster_A::> metrocluster switchover
```

Warning: negotiated switchover is about to start. It will stop all the data

Vservers on cluster "cluster_B" and
automatically re-start them on cluster
"cluster_A". It will finally gracefully shutdown
cluster "cluster_B".

Do you want to continue? {y|n}: y

7. Verify that all cluster SVMs are in a health state:

```
metrocluster vservers show
```

8. Resynchronize the data aggregates on the "surviving" cluster:

```
metrocluster heal -phase aggregates
```

After upgrading MetroCluster IP configurations to ONTAP 9.5 or later, the aggregates will be in a degraded state for a short period before resynchronizing and returning to a mirrored state.

```
cluster_A::> metrocluster heal -phase aggregates  
[Job 130] Job succeeded: Heal Aggregates is successful.
```

9. Verify that the healing operation was completed successfully:

```
metrocluster operation show
```

```
cluster_A::> metrocluster operation show  
Operation: heal-aggregates  
State: successful  
Start Time: MM/DD/YYYY TIME  
End Time: MM/DD/YYYY TIME  
Errors: -
```

10. Resynchronize the root aggregates on the "surviving" cluster:

```
metrocluster heal -phase root-aggregates
```

```
cluster_A::> metrocluster heal -phase root-aggregates  
[Job 131] Job succeeded: Heal Root Aggregates is successful.
```

11. Verify that the healing operation was completed successfully:

```
metrocluster operation show
```

```
cluster_A::> metrocluster operation show  
Operation: heal-root-aggregates  
State: successful  
Start Time: MM/DD/YYYY TIME  
End Time: MM/DD/YYYY TIME  
Errors: -
```

12. On the halted cluster, boot the node from the LOADER prompt:

```
boot_ontap
```

13. Wait for the boot process to finish, and then verify that all cluster SVMs are in a health state:

```
metrocluster vservers show
```

14. Perform a switchback from the “surviving” cluster:

```
metrocluster switchback
```

15. Verify that the switchback was completed successfully:

```
metrocluster operation show
```

```
cluster_A::> metrocluster operation show  
Operation: switchback  
State: successful  
Start Time: MM/DD/YYYY TIME  
End Time: MM/DD/YYYY TIME  
Errors: -
```

16. Verify that all cluster SVMs are in a health state:

```
metrocluster vserver show
```

17. Repeat all previous steps on the other cluster.
18. Verify that the MetroCluster configuration is healthy:
 - a. Check the configuration:

```
metrocluster check run
```

```
cluster_A::> metrocluster check run
Last Checked On: MM/DD/YYYY TIME
Component          Result
-----
nodes               ok
lifs                ok
config-replication ok
aggregates          ok
4 entries were displayed.
```

Command completed. Use the "metrocluster check show -instance" command or sub-commands in "metrocluster check" directory for detailed results.

To check if the nodes are ready to do a switchover or switchback operation, run "metrocluster switchover -simulate" or "metrocluster switchback -simulate", respectively.

- b. If you want to view more detailed results, use the metrocluster check run command:

```
metrocluster check aggregate show
```

```
metrocluster check config-replication show
```

```
metrocluster check lif show
```

```
metrocluster check node show
```

- c. Set the privilege level to advanced:

```
set -privilege advanced
```

d. Simulate the switchover operation:

```
metrocluster switchover -simulate
```

e. Review the results of the switchover simulation:

```
metrocluster operation show
```

```
cluster_A::*> metrocluster operation show
Operation: switchover
State: successful
Start time: MM/DD/YYYY TIME
End time: MM/DD/YYYY TIME
Errors: -
```

f. Return to the admin privilege level:

```
set -privilege admin
```

g. Repeat these substeps on the other cluster.

After you finish

Perform any [post-upgrade tasks](#).

Related information

[MetroCluster Disaster recovery](#)

Manual disruptive ONTAP upgrade using the CLI

If you can take your cluster offline to upgrade to a new ONTAP release, then you can use the disruptive upgrade method. This method has several steps: disabling storage failover for each HA pair, rebooting each node in the cluster, and then reenabling storage failover.

- You must [download](#) and [install](#) the software image.
- If you are operating in a SAN environment, all SAN clients must be shut down or suspended until the upgrade is complete.

If SAN clients are not shut down or suspended prior to a disruptive upgrade, then the client file systems and applications suffer errors that might require manual recovery after the upgrade is completed.

In a disruptive upgrade, downtime is required because storage failover is disabled for each HA pair, and each

node is updated. When storage failover is disabled, each node behaves as a single-node cluster; that is, system services associated with the node are interrupted for as long as it takes the system to reboot.

Steps

- 1. Set the privilege level from admin 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 * -iscurrent false} -isdefault true
```

This command uses an extended query to change the target ONTAP software image (which is installed as the alternate image) to be the default image for each node.

- 3. Verify that the new ONTAP software image is set as the default image:

```
system image show
```

In the following example, image 2 is the new ONTAP version and is set as the default image on both nodes:

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

- 4. Perform either one of the following steps:

| If the cluster consists of... | Do this... |
|-------------------------------|----------------------------|
| One node | Continue to the next step. |

| If the cluster consists of... | Do this... |
|-------------------------------|--|
| Two nodes | <p>a. Disable cluster high availability:</p> <pre>cluster ha modify -configured false</pre> <p>Enter y to continue when prompted.</p> <p>b. Disable storage failover for the HA pair:</p> <pre>storage failover modify -node * -enabled false</pre> |
| More than two nodes | <p>Disable storage failover for each HA pair in the cluster:</p> <pre>storage failover modify -node * -enabled false</pre> |

5. Reboot a node in the cluster:

```
system node reboot -node nodename -ignore-quorum-warnings
```



Do not reboot more than one node at a time.

The node boots the new ONTAP image. The ONTAP login prompt appears, indicating that the reboot process is complete.

6. After the node or set of nodes has rebooted with the new ONTAP image, set the privilege level to advanced:

```
set -privilege advanced
```

Enter **y** when prompted to continue

7. Confirm that the new software is running:

```
system node image show
```

In the following example, image1 is the new ONTAP version and is set as the current version on node0:

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

| Node | Image | Is Default | Is Current | Version | Install Date |
|-------|--------|---------------|---------------|---------|-----------------|
| node0 | image1 | true | true | X.X.X | MM/DD/YYYY TIME |
| | image2 | false | false | Y.Y.Y | MM/DD/YYYY TIME |
| node1 | image1 | true | false | X.X.X | MM/DD/YYYY TIME |
| | image2 | false | true | Y.Y.Y | MM/DD/YYYY TIME |

4 entries were displayed.

8. Verify that the upgrade is completed successfully:

- a. Set the privilege level to advanced:

```
set -privilege advanced
```

- b. Verify that the upgrade status is complete for each node:

```
system node upgrade-revert show -node nodename
```

The status should be listed as complete.

If the status is not complete, [contact NetApp Support](#) immediately.

- c. Return to the admin privilege level:

```
set -privilege admin
```

9. Repeat Steps 2 through 8 for each additional node.

10. If the cluster consists of two or more nodes, enable storage failover for each HA pair in the cluster:

```
storage failover modify -node * -enabled true
```

11. If the cluster consists of only two nodes, enable cluster high availability:

```
cluster ha modify -configured true
```


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.