



ONTAP® 9

Upgrade and Revert/Downgrade Guide

July 2019 | 215-11143_NO
doccomments@netapp.com

Updated for ONTAP 9.6

 **NetApp®**

Contents

Deciding whether to use this guide	5
When to revert and when to call technical support	5
Selecting your upgrade, downgrade, or revert procedure	7
Updating software on ONTAP clusters	8
Cluster software update workflow	8
Planning your update	9
Planning your update without Upgrade Advisor	9
Downgrade process considerations	15
Preparing to update the cluster	15
Creating a performance baseline with Perfstat Converged	16
Verifying that the cluster is ready	16
Preparing the ONTAP software for the update	25
Obtaining ONTAP software images	31
Installing the ONTAP software image	32
Selecting your update method for non-MetroCluster configurations	33
Selecting your update method for MetroCluster configurations	35
Upgrading an ONTAP cluster using the automated method	36
Requesting notification of issues encountered in nondisruptive upgrades	37
Performing an automatic nondisruptive upgrade using the CLI	38
Resuming an upgrade after an error in the automated upgrade process	40
Upgrading or downgrading a cluster nondisruptively by using the rolling upgrade method	40
Updating the first node in an HA pair	41
Updating the partner node in an HA pair	45
Updating a MetroCluster configuration using the manual method	50
Downgrade requirements for MetroCluster configurations	50
Updating a four- or eight-node MetroCluster configuration manually	51
Updating a two-node MetroCluster configuration in ONTAP 9.2 or earlier	65
Downgrading a two-node MetroCluster configuration disruptively	68
Updating an ONTAP cluster disruptively	69
Performing an automated upgrade on a single-node cluster	71
Completing post-upgrade or downgrade tasks for the cluster	73
Verifying the cluster version	74
Verifying cluster health (verifying storage health)	74
Verifying storage health (completing post-upgrade or downgrade tasks)	75
Verifying networking and storage status for MetroCluster configurations (post-upgrade or downgrade)	76
Verifying the SAN configuration after an upgrade	78
Reconfiguring KMIP server connections after upgrading to ONTAP 9.3 or later	78

Enabling and reverting LIFs to home ports (post-upgrade or downgrade tasks for the cluster)	79
Relocating moved load-sharing mirror source volumes	80
Resuming SnapMirror operations	80
Setting the desired NT ACL permissions display level for NFS clients	81
Enforcing SHA-2 on administrator account passwords	82
When you need to update the Disk Qualification Package	83
Reverting clusters to an earlier ONTAP release	84
When to revert and when to call technical support	84
Cluster revert workflow	84
Planning your reversion	85
Reviewing pre-reversion resources	85
Reviewing cluster reversion requirements	86
Preparing to revert ONTAP clusters	87
Verifying that the cluster is ready to be reverted	88
Preparing to revert production clusters	91
Obtaining ONTAP software images	97
Reverting an ONTAP cluster	99
Completing post-reversion tasks	102
Enabling automatic switchover for MetroCluster configurations	102
Verifying cluster health (completing post-reversion tasks)	103
Verifying storage health (completing post-reversion tasks)	104
Enabling and reverting LIFs to home ports (completing post-reversion tasks)	105
Preparing Snapshot copies after reverting	106
Verifying client access (CIFS and NFS)	106
Verifying IPv6 firewall entries	107
Reverting password hash function to the supported encryption type	108
Considerations for whether to manually update the SP firmware	108
Optimal service availability during upgrades	109
Considerations for services and protocols during upgrades	109
Considerations for stateless protocols	109
Considerations for session-oriented protocols	110
How firmware is updated during the ONTAP upgrade	110
Understanding background disk firmware updates	111
Copyright	112
Trademark	113
How to send comments about documentation and receive update notifications	114

Deciding whether to use the Upgrade and Revert/Downgrade Guide

This guide describes how to manually upgrade, downgrade, or revert an ONTAP cluster or a MetroCluster configuration using the manual nondisruptive or disruptive upgrade process. It also describes how to perform an automated upgrade using the command line interface (CLI).

Beginning in ONTAP 9.3, the automated upgrade procedure is the preferred upgrade method for all configurations, including MetroCluster configurations. Only patch updates are supported for automated updates on MetroCluster configurations prior to ONTAP 9.3.

If you prefer to a user interface driven upgrade instead of the CLI, use ONTAP System Manager, See the *Upgrade Express Guide*.

Software express upgrade

You should only use the manual upgrade procedures if you require the level of control and monitoring that the manual procedure provides.

Unless otherwise indicated, the requirements and procedures in this guide apply to all platforms supported in ONTAP 9 and to the upgrade and revert/downgrade paths outlined in Cluster update requirements.

Related concepts

Cluster update requirements on page 10

When to revert and when to call technical support

You can downgrade or revert without assistance when downgrading or reverting new or test clusters, but you should call technical support if you encounter problems during or after upgrade, or if you want to downgrade or revert a production cluster.

You can revert or downgrade to an allowed ONTAP release without assistance from technical support only in the following scenarios:

- You upgraded to a new release on a test cluster and you want to return to the original release when testing is completed.
- You are configuring a new cluster—running a later release of ONTAP and not yet in production—in an environment in which you have standardized on an earlier ONTAP release.

If the upgrade fails, *do not* attempt to revert ONTAP in a production environment without assistance. If you encounter any of the following circumstances, contact technical support immediately:

- The upgrade process fails and cannot finish.
- The upgrade process finishes, but the cluster is unusable in a production environment.
- The upgrade process finishes and the cluster goes into production, but you are not satisfied with its behavior.
- The upgrade process finishes for some but not all of the nodes, and you decide that you want to revert.

If you created volumes in ONTAP 9.5 or later and you need to revert to an earlier version, contact technical support to confirm if any of the volumes use adaptive compression. Volumes using adaptive compression must be uncompressed before reverting.

6 | Upgrade and Revert/Downgrade Guide

Related concepts

Cluster update requirements on page 10

Selecting your upgrade, downgrade, or revert procedure

When you update the ONTAP software, you must use a different procedure depending on whether you are upgrading, downgrading, or reverting the software.

An ONTAP software *update* involves one of the following possible activities.

Software upgrade

The software version is changed from an earlier version to a later version of ONTAP. For example, from ONTAP 9.3 to ONTAP 9.4.

[Updating software on ONTAP clusters](#) on page 8

Software downgrade or Software revert

The software version is changed from a later version to an earlier version of ONTAP. For example, from ONTAP 9.3 to ONTAP 9.2. Verify your cluster requirements to determine if you need to downgrade or revert.

[Updating software on ONTAP clusters](#) on page 8

[Reverting clusters to an earlier ONTAP release](#) on page 84

Updating software on ONTAP clusters

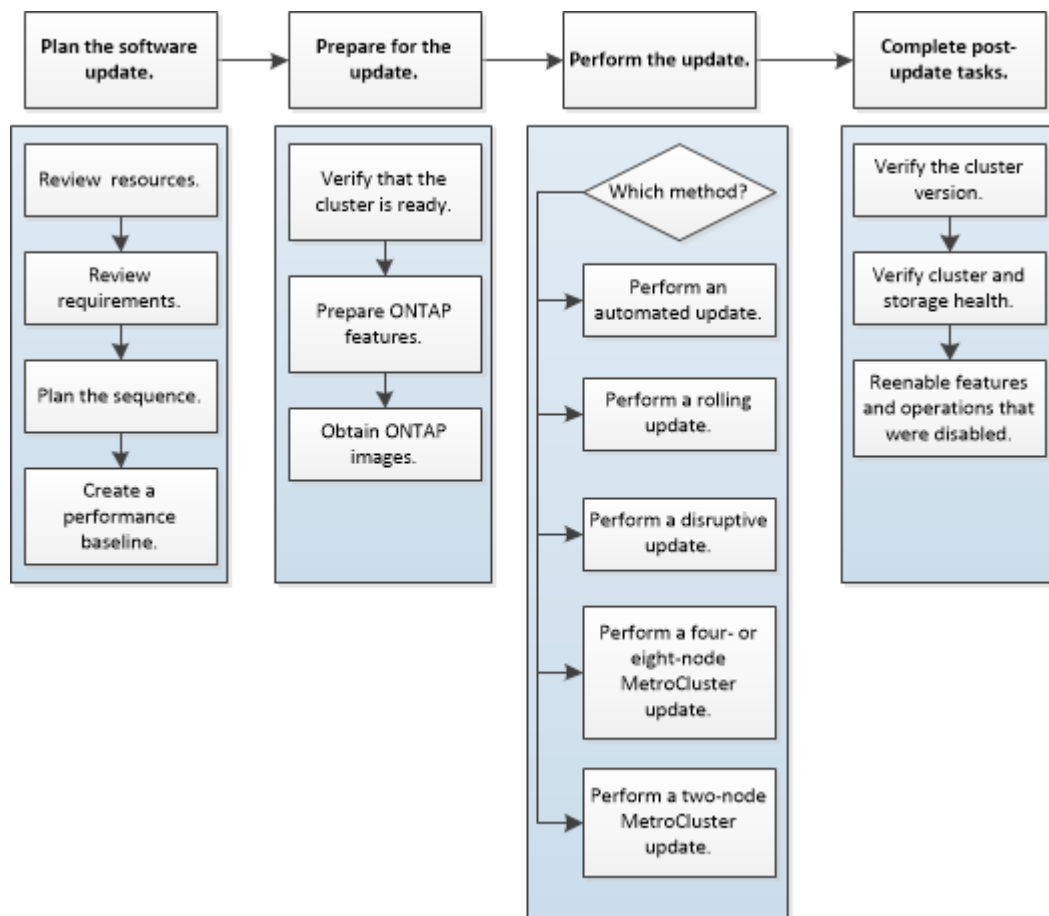
Upgrading or downgrading a cluster to the current ONTAP release requires planning, preparation, the upgrade or downgrade itself, and several post-upgrade or downgrade procedures.

The software update process includes the following phases:

- Planning for the update
- Preparing for the update
- Performing the update
- Completing post-update tasks

Cluster software update workflow

You can use the cluster software update workflow to perform the entire process.



Planning your update

It is a best practice to use Upgrade Advisor in Active IQ to plan your upgrade. If you cannot use Upgrade Advisor, you should create your own upgrade plan manually by using guidelines provided in this guide.

Related concepts

[Updating software on ONTAP clusters](#) on page 8

[Reverting clusters to an earlier ONTAP release](#) on page 84

Related tasks

[Planning your update without Upgrade Advisor](#) on page 9

If you are not using Upgrade Advisor, you must manually determine your plan for the update operation.

Related information

[NetApp Active IQ](#)

Planning your update without Upgrade Advisor

If you are not using Upgrade Advisor, you must manually determine your plan for the update operation.

Steps

1. [Reviewing pre-update resources](#) on page 9
2. [Reviewing cluster upgrade/downgrade requirements](#) on page 10
3. [Verifying cluster upgrade limits](#) on page 14

Reviewing pre-update resources

Before updating the ONTAP software, you should review resources to understand issues you must resolve, understand new system behavior in the target release, and confirm hardware support.

Steps

1. Review the *Release Notes* for the target release.

[ONTAP 9 Release Notes](#)

The “Important cautions” section describes potential issues that you should be aware of before upgrading to the new release. The “New and changed features” and “Known problems and limitations” sections describe new system behavior after upgrading to the new release.

2. Confirm that your hardware platform is supported in the target release.

[NetApp Hardware Universe](#)

3. Confirm that your cluster and management switches are supported in the target release.

Your NX-OS (cluster network switches), IOS (management network switches), and reference configuration file (RCF) software versions must be compatible with the version of ONTAP to which you are upgrading.

[NetApp Interoperability Matrix Tool](#)

4. If your cluster and management switches do not have the minimum software versions for the target ONTAP release, upgrade to supported software versions.

NetApp Downloads: Cisco Ethernet Switch

NetApp Downloads: NetApp Ethernet Switch

5. If your cluster is configured for SAN, confirm that the SAN configuration is fully supported.

All SAN components—including the target ONTAP software version, host OS and patches, required Host Utilities software, multipathing software, and adapter drivers and firmware—should be supported.

NetApp Interoperability Matrix Tool

6. If you are transitioning from 7-Mode using the 7-Mode Transition Tool, confirm that the tool supports transition to the ONTAP version to which you are upgrading.

All the projects in the tool must be in the completed or aborted state before you upgrade the 7-Mode Transition Tool that supports the ONTAP version to which you are upgrading.

7-Mode Transition Tool installation and administration

Reviewing cluster upgrade/downgrade requirements

Before updating the ONTAP software, you must verify that your cluster meets the general requirements. Some configurations and features also have requirements that you should understand.

Cluster update requirements

There are release and configuration requirements that your cluster should meet before you perform an update. Additionally, there are mixed version requirements that you should be aware of while you are performing the update.

Release requirements

The version of ONTAP that you can upgrade or downgrade to varies based on the version of ONTAP currently running on your nodes. You can determine the current version of ONTAP running on each node by using the `system image show` command.

You can upgrade from...	To...
ONTAP 9.5	ONTAP 9.6
ONTAP 9.4	ONTAP 9.5
ONTAP 9.3	ONTAP 9.4 or ONTAP 9.5
ONTAP 9.2	ONTAP 9.3
ONTAP 9.1	ONTAP 9.2 or 9.3 Note: If you are running a release earlier than ONTAP 9.1, you cannot upgrade directly to ONTAP 9.2 or ONTAP 9.3. You must upgrade to ONTAP 9.1 first, then upgrade to ONTAP 9.2 or ONTAP 9.3.
ONTAP 9	ONTAP 9.1

You can upgrade from...	To...
Data ONTAP 8.3.x	ONTAP 9 or 9.1 Note: If you are running a release earlier than Data ONTAP 8.3.x, you cannot upgrade directly to ONTAP 9 or 9.1. You must upgrade to Data ONTAP 8.3.x first, then upgrade to ONTAP 9 or 9.1.
Data ONTAP 8.2.x	Data ONTAP 8.3.x

You can downgrade from...	To...
ONTAP 9.1	ONTAP 9 Note: Downgrade of a two-node MetroCluster configuration from ONTAP 9.1 to 9 is disruptive.

You must perform a revert from...	To...
ONTAP 9.6	ONTAP 9.5
ONTAP 9.5	ONTAP 9.4
ONTAP 9.4	ONTAP 9.3
ONTAP 9.3	ONTAP 9.2
ONTAP 9.2	ONTAP 9.1
ONTAP 9.1 or ONTAP 9	Data ONTAP 8.3.x

Mixed version requirements

Beginning with ONTAP 9.3, by default, you cannot join new nodes to the cluster that are running a version of ONTAP that is different from the version running on the existing nodes. If you plan to add new nodes to your cluster that are running a version of ONTAP that is later than the nodes in your existing cluster, you should upgrade the nodes in your cluster to the later version first, then add the new nodes.

Mixed version clusters are not recommended, but in certain cases you might need to temporarily enter a mixed version state. For example, you need to enter a mixed version state if you are upgrading to a later version of ONTAP that is not supported on certain nodes in your existing cluster. In this case, you should upgrade the nodes that do support the later version of ONTAP, then unjoin the nodes that do not support the version of ONTAP you are upgrading to using the advance privilege `cluster unjoin -skip-lastlow-version-node check` command.

You might also need to enter a mixed version state for a technical refresh or an interrupted upgrade. In such cases you can override the ONTAP 9.3 default behavior and join nodes of a different version using the following advance privilege commands:

- `cluster join -allow-mixed-version-join`
- `cluster add-node -allow-mixed-version-join`

When you have to enter a mixed version state, you should complete the upgrade as quickly as possible. An HA pair must not run an ONTAP version from a release that is different from other HA pairs in the cluster for more than seven days. For correct cluster operation, the period the cluster is in a mixed version state should be as short as possible.

When the cluster is in a mixed version state, you should not enter any commands that alter the cluster operation or configuration except as necessary to satisfy the upgrade requirements.

Guidelines for estimating the duration of the upgrade process

You should plan for at least 30 minutes to complete preparatory steps, 60 minutes to upgrade each HA pair, and at least 30 minutes to complete post-upgrade steps.

The upgrade duration guidelines are based on typical configurations and workloads. You can use these guidelines to estimate the time it will take to perform a nondisruptive upgrade in your environment. However, the actual duration of your upgrade process will depend on your individual environment and the number of nodes.

Upgrade considerations for SVM routing

The routing table for an SVM determines the network path the SVM uses to communicate with a destination. It's important to understand how routing tables work so that you can prevent network problems before they occur.

Routing rules are as follows:

- ONTAP routes traffic over the most specific available route.
- ONTAP routes traffic over a default gateway route (having 0 bits of netmask) as a last resort, when more specific routes are not available.

In the case of routes with the same destination, netmask, and metric, there is no guarantee that the system will use the same route after a reboot or after an upgrade. This is especially an issue if you have configured multiple default routes.

It is a best practice to configure one default route for an SVM. To avoid disruption, you should ensure that the default route is able to reach any network address that is not reachable by a more specific route. For more information, see [NetApp KB Article 1000317: Network access might be disrupted by incorrect routing configuration in clustered Data ONTAP](#).

Upgrade considerations for root-data partitioning and root-data-data partitioning

Root-data partitioning and root-data-data-partitioning is supported for some platform models and configurations. This partitioning capability is enabled during system initialization; it cannot be applied to existing aggregates.

For information about migrating your data to a node that is configured for root-data partitioning or root-data-data partitioning, contact your account team or partner organization.

Related information

[ONTAP concepts](#)

Upgrade requirements for SnapMirror

You must perform certain tasks to successfully upgrade a cluster that is running SnapMirror.

- If you are upgrading clusters with an inter-cluster DP SnapMirror relationship, you must upgrade the destination cluster before you upgrade the source cluster.
- Before upgrading a cluster that is running SnapMirror, SnapMirror operations must be suspended for each node that contains destination volumes, and each peered SVM must have a unique name across the clusters.

For SnapMirror volume replication, the destination node must use an ONTAP version that is equal to or later than that of the SnapMirror source node. To prevent SnapMirror transfers from failing, you must suspend SnapMirror operations and, in some cases, upgrade destination nodes before upgrading source nodes. The following table describes the two options for suspending SnapMirror operations.

Option	Description	Upgrade destination nodes before source nodes?
Suspend SnapMirror operations for the duration of the NDU (nondisruptive upgrade).	The simplest method for upgrading in a SnapMirror environment is to suspend all SnapMirror operations, perform the upgrade, and then resume the SnapMirror operations. However, no SnapMirror transfers will occur during the entire NDU. You must use this method if your cluster contains nodes that are mirroring volumes to each other.	No, the nodes can be upgraded in any order.
Suspend SnapMirror operations one destination volume at a time.	You can suspend SnapMirror transfers for a particular destination volume, upgrade the node (or HA pair) that contains the destination volume, upgrade the node (or HA pair) that contains the source volume, and then resume the SnapMirror transfers for the destination volume. By using this method, SnapMirror transfers for all other destination volumes can continue while the nodes that contain the original destination and source volumes are upgraded.	Yes.

SVM peering requires SVM names to be unique across clusters. You should name SVMs with a unique fully qualified domain name (FQDN), for example, “dataVserver.HQ” or “mirrorVserver.Offsite”. Using the FQDN naming style makes it much easier to make sure of uniqueness.

Related information

[ONTAP concepts](#)

Upgrade requirements for MetroCluster configurations

If you have to upgrade a MetroCluster configuration, you should be aware of some important requirements.

Required methods for performing major and minor upgrades of MetroCluster configurations

Patch upgrades to MetroCluster configurations can be performed with automatic non-disruptive upgrade (NDU) procedure.

Starting with ONTAP 9.3, major upgrades to MetroCluster configurations can be performed with automatic non-disruptive upgrade (NDU) procedure. On systems running ONTAP 9.2 or earlier, major upgrades to MetroCluster configurations must be performed with the NDU procedure that is specific to MetroCluster configurations.

General requirements

- Both clusters must be running the same version of ONTAP.
You can verify the ONTAP version by using the `version` command.

- The MetroCluster configuration must be in either normal or switchover mode.
- For all configurations except two-node clusters, you can nondisruptively upgrade both clusters at the same time.
For nondisruptive upgrade in two-node clusters, the clusters must be upgraded one node at a time.
- The aggregates in both clusters must not be in `resyncing RAID` status.
During MetroCluster healing, the mirrored aggregates are resynchronized. You can verify whether the MetroCluster configuration is in this state by using the `storage aggregate plex show -in-progress true` command. If any aggregates are being synchronized, you should not perform an upgrade until the resynchronization is complete.
- Negotiated switchover operations will fail while the upgrade is in progress.
To avoid issues with upgrade or revert operations, do not attempt an unplanned switchover during an upgrade or revert operation unless all nodes on both clusters are running the same version of ONTAP.

Configuration requirements for normal operation

- The source SVM LIFs must be up and located on their home nodes.
Data LIFs for the destination SVMs are not required to be up or to be on their home nodes.
- All aggregates at the local site must be online.
- All root and data volumes owned by the local cluster's SVMs must be online.

Configuration requirements for switchover

- All LIFs must be up and located on their home nodes.
- All aggregates must be online, except for the root aggregates at the DR site.
Root aggregates at the DR site are offline during certain phases of switchover.
- All volumes must be online.

Related tasks

[Verifying networking and storage status for MetroCluster configurations \(cluster is ready\)](#) on page 21

Upgrade considerations for SnapLock

SnapLock does not allow the download of certain kernel versions if these are qualified as bad SnapLock releases or if SnapLock is disabled in those releases. These download restrictions only apply if the node has SnapLock data.

Verifying cluster upgrade limits

Before upgrading the ONTAP software, you must verify that your cluster does not exceed the platform system limits. SAN also has limits that you should verify in addition to the platform system limits.

Steps

1. Verify that the cluster does not exceed the system limits for your platform.
[NetApp Hardware Universe](#)
2. If your cluster is configured for SAN, verify that it does not exceed the configuration limits for FC, FCoE, and iSCSI.
[SAN configuration](#)

3. Determine the CPU and disk utilization:

```
node run -node node_name -command sysstat -c 10 -x 3
```

You should monitor CPU and disk utilization for 30 seconds. The values in the `CPU` and `Disk Util` columns should not exceed 50% for all 10 measurements reported. No additional load should be added to the cluster until the upgrade is complete.

Downgrade process considerations

You need to know about downgrade issues and limitations before downgrading clusters to an earlier version of ONTAP.

You should be aware of the following:

- You can only downgrade from ONTAP 9.1 to 9.0. For all other versions of ONTAP 9, you must perform a revert.
- If the version of ONTAP you are downgrading to has a different BIOS version than your current ONTAP version, you should downgrade your BIOS before you downgrade ONTAP.
- Downgrading affects all nodes in the cluster.
- You can downgrade ONTAP nondisruptively, except for single-node clusters, which lack hardware redundancy.
During the downgrade process, the cluster remains online and continues to serve data.
- If your cluster serves CIFS clients, nondisruptive downgrades are supported for Hyper-V and SQL Server over SMB solutions.
These solutions enable the application servers and the contained virtual machines or databases to stay online and to provide continuous availability during the ONTAP downgrade.
For all other CIFS configurations, client sessions are terminated. You should direct users to end their sessions before you downgrade to prevent data loss.
- ONTAP clusters can operate for a limited time in a *mixed version* state, in which nodes in a cluster are running different versions of ONTAP; however, the update is not complete until all nodes are running the new target release.
When the cluster is in a mixed version state, you should not enter any commands that alter the cluster operation or configuration except as necessary to satisfy upgrade requirements. You should complete the update as quickly as possible; do not allow the cluster to remain in a mixed version state longer than necessary. An HA pair must not run an ONTAP version from a release that is different from other HA pairs in the cluster for more than seven days.

Related information

[SMB/CIFS management](#)

Preparing to update the cluster

Before performing an upgrade or downgrade you must manually check that the cluster is ready, make any required configuration changes, and obtain and install the target ONTAP images.

Steps

1. [Creating a performance baseline with Perfstat Converged](#) on page 16
The Performance and Statistics Collector (Perfstat Converged) is a cluster diagnostics data collection tool, available on the NetApp Support Site that enables you to establish a performance baseline for comparison after the upgrade. You should create a Perfstat report before upgrading.
2. [Verifying that the cluster is ready](#) on page 16

Before you perform the upgrade or downgrade, you should verify that your cluster configuration is healthy.

3. [Preparing the ONTAP software for the update](#) on page 25
Some ONTAP features have configuration requirements that must be completed before the cluster software version can be updated.
4. [Obtaining ONTAP software images](#) on page 31
For ONTAP 9.4 and later, you can copy the ONTAP software image from the NetApp Support Site to a local folder. For upgrades from ONTAP 9.3 or earlier, you must copy the ONTAP software image to an HTTP server or FTP server on your network.
5. [Installing the ONTAP software image](#) on page 32
You must install the target software image on the cluster's nodes.

Creating a performance baseline with Perfstat Converged

The Performance and Statistics Collector (Perfstat Converged) is a cluster diagnostics data collection tool, available on the NetApp Support Site that enables you to establish a performance baseline for comparison after the upgrade. You should create a Perfstat report before upgrading.

Before you begin

The diag user account must be unlocked.

[System administration](#)

Steps

1. Download Perfstat Converged from the NetApp Support Site to create a Perfstat report during a typical usage time.

[NetApp Downloads: Performance and Statistics Collector \(Perfstat\)](#)

This takes about 30 minutes.

2. Enter the following command during a typical usage period:

```
perfstat8 cluster_management_IP_address -m c -t 4 -i 5 -z
```

After you finish

You should retain the output file for several weeks after the ONTAP upgrade is complete to compare with the performance of the new version.

Verifying that the cluster is ready

Before you perform the upgrade or downgrade, you should verify that your cluster configuration is healthy.

Checking for common configuration errors using Config Advisor

You can use the Config Advisor tool to check for common configuration errors.

About this task

Config Advisor is a configuration validation and health check tool for NetApp systems. This tool can be deployed at both secure sites and nonsecure sites for data collection and system analysis.

Note: Support for Config Advisor is limited and is available only online.

Steps

1. Log in to the NetApp Support Site, and then navigate to **Downloads > Software > ToolChest**.
[NetApp Downloads: Config Advisor](#)
2. Click **Config Advisor**.
3. Download, install, and run Config Advisor by following the directions on the web page.
4. After running Config Advisor, review the tool's output, and follow the recommendations that are provided to address any issues that are discovered by the tool.

Checking for MetroCluster configuration errors with Config Advisor

You can go to the NetApp Support Site and download the Config Advisor tool to check for common configuration errors.

About this task

Config Advisor is a configuration validation and health check tool. You can deploy it at both secure sites and non-secure sites for data collection and system analysis.

Note: Support for Config Advisor is limited, and available only online.

Steps

1. Go to the Config Advisor download page and download the tool.
[NetApp Downloads: Config Advisor](#)
2. Run Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

Verifying LDAP status

If LDAP is used by your storage virtual machines (SVMs), you must have an established LDAP connection to perform a nondisruptive upgrade. You should verify the LDAP connection before you begin the upgrade.

Steps

1. Check the LDAP status:
`ldap check -vserver vserver_name`
2. If the LDAP status is down, modify it:
`ldap client modify -client-config LDAP_client -ldap-servers ip_address`
3. Verify that the LDAP status is up:
`ldap check -vserver vserver_name`

Verifying DNS server status

Before and after performing a nondisruptive upgrade, you should verify the status of your Domain Name Service (DNS) server.

About this task**Steps**

1. Check the status of your DNS servers:

```
dns check -vserver vserver_name
```

An up status indicates the service is running. A down status indicates that the service is not running.

2. If the DNS server is down, modify it:

```
dns modify -vserver vserver_name -domains domain_name -name-servers
name_server_ipaddress
```

3. Verify the status of the DNS server is up.

Verifying HA status

Before performing a nondisruptive upgrade, you should verify that storage failover is enabled for each HA pair. If the cluster consists of only two nodes, you should also verify that cluster HA is enabled.

About this task

You do not need to verify the HA status if you plan to perform a disruptive upgrade, because this upgrade method does not require storage failover.

Steps

1. Verify that storage failover is enabled and possible for each HA pair:

```
storage failover show
```

Example

This example shows that storage failover is enabled and possible on node0 and node1:

```
cluster1::> storage failover show
Node           Partner           Takeover
Possible State
-----
node0          node1             true    Connected to node1
node1          node0             true    Connected to node0
2 entries were displayed.
```

If necessary, you can enable storage failover by using the `storage failover modify` command.

2. If the cluster consists of only two nodes (a single HA pair), verify that cluster HA is configured:

```
cluster ha show
```

Example

This example shows that cluster HA is configured:

```
cluster1::> cluster ha show
High Availability Configured: true
```

If necessary, you can enable cluster HA by using the `cluster ha modify` command.

Verifying cluster health (verifying that the cluster is ready)

Before and after you upgrade, revert, or downgrade a cluster, you should verify that the nodes are healthy and eligible to participate in the cluster, and that the cluster is in quorum.

Steps

1. Verify that the nodes in the cluster are online and are eligible to participate in the cluster:

cluster show

Example

```
cluster1::> cluster show
Node                Health  Eligibility
-----
node0                true   true
node1                true   true
```

If any node is unhealthy or ineligible, check EMS logs for errors and take corrective action.

2. Set the privilege level to advanced:

```
set -privilege advanced
```

3. Enter **y** to continue.

4. Verify the configuration details for each RDB process.

- The relational database epoch and database epochs should match for each node.
- The per-ring quorum master should be the same for all nodes.
Note that each ring might have a different quorum master.

To display this RDB process...	Enter this command...
Management application	cluster ring show -unitname mgmt
Volume location database	cluster ring show -unitname vldb
Virtual-Interface manager	cluster ring show -unitname vifmgr
SAN management daemon	cluster ring show -unitname bcomd

Example

This example shows the volume location database process:

```
cluster1::*> cluster ring show -unitname vldb
Node      UnitName Epoch  DB Epoch DB Trnxs Master  Online
-----
node0     vldb     154    154     14847  node0  master
node1     vldb     154    154     14847  node0  secondary
node2     vldb     154    154     14847  node0  secondary
node3     vldb     154    154     14847  node0  secondary
4 entries were displayed.
```

5. If you are operating in a SAN environment, verify that each node is in a SAN quorum:

```
event log show -messagename scsiblade.*
```

The most recent `scsiblade` event message for each node should indicate that the scsi-blade is in quorum.

Example

```
cluster1::*> event log show -messagename scsiblade.*
Time                Node      Severity  Event
-----
MM/DD/YYYY TIME  node0     INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
MM/DD/YYYY TIME  node1     INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
```

- Return to the admin privilege level:

```
set -privilege admin
```

Related information

[System administration](#)

Verifying storage health (verifying that the cluster is ready)

Before and after you upgrade, revert, or downgrade a cluster, you should verify the status of your disks, aggregates, and volumes.

Steps

- If you are preparing to upgrade, revert, or downgrade, verify disk status:

To check for..	Do this..
Broken disks	<ol style="list-style-type: none"> Display any broken disks: <code>storage disk show -state broken</code> Remove or replace any broken disks.
Disks undergoing maintenance or reconstruction	<ol style="list-style-type: none"> Display any disks in maintenance, pending, or reconstructing states: <code>storage disk show -state maintenance pending reconstructing</code> Wait for the maintenance or reconstruction operation to finish before proceeding.

- Verify that all aggregates are online by displaying the state of physical and logical storage, including storage aggregates:

```
storage aggregate show -state !online
```

This command displays the aggregates that are *not* online. All aggregates must be online before and after performing a major upgrade or reversion.

Example

```
cluster1::> storage aggregate show -state !online
There are no entries matching your query.
```

- Verify that all volumes are online by displaying any volumes that are *not* online:

```
volume show -state !online
```

All volumes must be online before and after performing a major upgrade or reversion.

Example

```
cluster1::> volume show -state !online
There are no entries matching your query.
```

- Verify that there are no inconsistent volumes:

```
volume show -is-inconsistent true
```

If any inconsistent volumes are returned, you must contact NetApp Support before you proceed with the upgrade.

Related information[Logical storage management](#)**Verifying networking and storage status for MetroCluster configurations (cluster is ready)**

Before and after performing an update in a MetroCluster configuration, you should verify the status of the LIFs, aggregates, and volumes for each cluster.

Steps

1. Verify the LIF status:

network interface show

In normal operation, LIFs for source SVMs must have an admin status of up and be located on their home nodes. LIFs for destination SVMs are not required to be up or located on their home nodes. In switchover, all LIFs have an admin status of up, but they do not need to be located on their home nodes.

Example

```
cluster1::> network interface show
Vserver      Logical      Status      Network      Current      Current      Is
Interface    Admin/Oper   Address/Mask Node          Port         Home
-----
Cluster
  cluster1-a1_clus1
                up/up      192.0.2.1/24  cluster1-01  e2a         true
  cluster1-a1_clus2
                up/up      192.0.2.2/24  cluster1-01  e2b         true
cluster1-01
  clus_mgmt    up/up      198.51.100.1/24  cluster1-01  e3a         true
  cluster1-a1_inet4_intercluster1
                up/up      198.51.100.2/24  cluster1-01  e3c         true
  ...
27 entries were displayed.
```

2. Verify the state of the aggregates:

storage aggregate show -state !online

This command displays any aggregates that are *not* online. In normal operation, all aggregates located at the local site must be online. However, if the MetroCluster configuration is in switchover, root aggregates at the disaster recovery site are permitted to be offline.

Example

This example shows a cluster in normal operation:

```
cluster1::> storage aggregate show -state !online
There are no entries matching your query.
```

Example

This example shows a cluster in switchover, in which the root aggregates at the disaster recovery site are offline:

```
cluster1::> storage aggregate show -state !online
Aggregate      Size Available Used% State  #Vols  Nodes      RAID Status
-----
aggr0_b1
                0B      0B      0% offline  0  cluster2-01  raid_dp,
                mirror
                degraded
```

```

aggr0_b2          0B          0B      0% offline      0 cluster2-02      raid_dp,
mirror
degraded
2 entries were displayed.

```

3. Verify the state of the volumes:

```
volume show -state !online
```

This command displays any volumes that are *not* online.

If the MetroCluster configuration is in normal operation (it is not in switchover state), the output should show all volumes owned by the cluster's secondary SVMs (those with the SVM name appended with "-mc").

Those volumes come online only in the event of a switchover.

Example

This example shows a cluster in normal operation, in which the volumes at the disaster recovery site are not online.

```

cluster1::> volume show -state !online
(volume show)
Vserver  Volume      Aggregate  State  Type  Size  Available  Used%
-----
vs2-mc   vol1        aggr1_b1   -      RW    -      -          -
vs2-mc   root_vs2    aggr0_b1   -      RW    -      -          -
vs2-mc   vol2        aggr1_b1   -      RW    -      -          -
vs2-mc   vol3        aggr1_b1   -      RW    -      -          -
vs2-mc   vol4        aggr1_b1   -      RW    -      -          -
5 entries were displayed.

```

4. Verify that there are no inconsistent volumes:

```
volume show -is-inconsistent true
```

If any inconsistent volumes are returned, you must contact NetApp Support before you precede with the upgrade.

Related concepts

[Upgrade requirements for MetroCluster configurations](#) on page 13

Verifying that deduplicated volumes and aggregates contain sufficient free space

Before upgrading ONTAP, you must verify that any deduplicated volumes and the aggregates that contain them have sufficient free space for the deduplication metadata. If there is insufficient free space, deduplication will be disabled when the ONTAP upgrade is completed.

About this task

Each deduplicated volume must contain at least 4% free space. Each aggregate that contains a deduplicated volume must contain at least 3% free space.

Steps

1. Determine which volumes are deduplicated:

```
volume efficiency show
```

2. Determine the free space available on each volume that you identified:

```
df -vserver Vserver_name -volume volume_name
```

Each deduplicated volume must not contain more than 96% used capacity. If necessary, you can increase the sizes of any volumes that exceed this capacity.

[Logical storage management](#)

Example

In this example, the `capacity` field displays the percentage of used space on the deduplicated volume identified earlier (`vol_2`):

```
cluster1::> df -vserver vs2 -volume vol_2
Filesystem          kbytes    used    avail capacity  Mounted on
/vol/vol_2/        19456000  264000  19192000    1% /
/vol/vol_2/.snapshot 1024      0      1024      0% // .snapshot
2 entries were displayed.
```

3. Identify the free space available on each aggregate that contains a deduplicated volume:

```
df -A -aggregate aggregate_name
```

Each aggregate must not contain more than 97% used capacity. If necessary, you can increase the sizes of any aggregates that exceed this capacity.

*Disk and aggregate management***Example**

In this example, the `capacity` field displays the percentage of used space on the aggregate containing the deduplicated volume (`aggr_2`):

```
cluster1::> df -A -aggregate aggr_2
Aggregate          kbytes    used    avail capacity
aggr_2             344220000  20944000  323276000    6%
aggr_2/.snapshot   0          0          0          0%
2 entries were displayed.
```

Verifying the LIF failover configuration

Before you perform an upgrade, you must verify that the failover policies and failover groups are configured correctly.

Steps

1. Display the failover policy for each data LIF:

```
network interface show -role data -failover
```

Example

This example shows the default failover configuration for a two-node cluster with two data LIFs:

```
cluster1::> network interface show -role data -failover
Vserver   Logical   Home      Failover   Failover
Interface Interface Node:Port Policy      Group
-----
vs0
lif0      node0:e0b nextavail  system-defined
Failover Targets: node0:e0b, node0:e0c,
node0:e0d, node0:e0e,
node0:e0f, node1:e0b,
node1:e0c, node1:e0d,
node1:e0e, node1:e0f
vs1
lif1      node1:e0b nextavail  system-defined
Failover Targets: node1:e0b, node1:e0c,
node1:e0d, node1:e0e,
node1:e0f, node0:e0b,
node0:e0c, node0:e0d,
node0:e0e, node0:e0f
```

The `Failover Targets` field shows a prioritized list of failover targets for each LIF. For example, if `lif0` fails over from its home port (`e0b` on `node0`), it first attempts to fail over to

port `e0c` on `node0`. If `lif0` cannot fail over to `e0c`, it next attempts to fail over to port `e0d` on `node0`, and so on.

2. If you have LIFs on multiple IP subnets, verify that each LIF belongs to a failover group that contains ports on the same layer 2 broadcast domain.
A user-defined failover group must be configured for each VLAN or broadcast domain, and each LIF must subscribe to the corresponding failover group.
3. If the failover policy is set to `disabled` for any of the LIFs, use the `network interface modify` command to enable failover.
4. For each LIF, verify that the `Failover Targets` field includes data ports from a different node that will remain up while the LIF's home node is being upgraded.

You can use the `network interface failover-groups create` command to add a failover target to the failover group.

Related information

[Network and LIF management](#)

Ensuring that no jobs are running

Before updating or downgrading the ONTAP software, you must verify the status of cluster jobs. If any aggregate, volume, NDMP (dump or restore), or Snapshot jobs (such as create, delete, move, modify, replicate, and mount jobs) are running or queued, you must allow the jobs to finish successfully or stop the queued entries.

Steps

1. Review the list of any running or queued aggregate, volume, or Snapshot jobs:

```
job show
```

Example

```
cluster1::> job show
Job ID Name                Owing      Node      State
-----
8629  Vol Reaper                cluster1   -         Queued
      Description: Vol Reaper Job
8630  Certificate Expiry Check  cluster1   -         Queued
      Description: Certificate Expiry Check
.
.
.
```

2. Delete any running or queued aggregate, volume, or Snapshot copy jobs:

```
job delete -id job_id
```

Example

```
cluster1::> job delete -id 8629
```

3. Verify that no aggregate, volume, or Snapshot jobs are running or queued:

```
job show
```

Example

In this example, all running and queued jobs have been deleted:


```

cluster1::> job show
Job ID Name                Owing      Node      State
-----
9944  SnapMirrorDaemon_7_2147484678
      cluster1  node1     Dormant
      Description: Snapmirror Daemon for 7_2147484678
18377 SnapMirror Service Job
      cluster1  node0     Dormant
      Description: SnapMirror Service Job
2 entries were displayed

```

Verifying the SAN configuration

Upgrading in a SAN environment changes which paths are direct. Therefore, before performing an upgrade, you should verify that each host is configured with the correct number of direct and indirect paths, and that each host is connected to the correct LIFs.

Steps

1. On each host, verify that a sufficient number of direct and indirect paths are configured, and that each path is active.

Each host must have a path to each node in the cluster.

2. Verify that each host is connected to a LIF on each node.

You should record the list of initiators for comparison after the upgrade.

For..	Enter..
iSCSI	<code>iscsi initiator show -fields igroup,initiator-name,tpgroup</code>
FC	<code>fcp initiator show -fields igroup,wwpn,lif</code>

Preparing the ONTAP software for the update

Some ONTAP features have configuration requirements that must be completed before the cluster software version can be updated.

Verifying that the netgroup file is present on all nodes

If you have loaded netgroups into storage virtual machines (SVMs), before you upgrade or revert, you must verify that the netgroup file is present on each node. A missing netgroup file on a node can cause an upgrade or revert to fail.

About this task

The *NFS Reference* contains more information about netgroups and loading them from a URI.

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Display the netgroup status for each SVM:

```
vserver services netgroup status
```

3. Verify that for each SVM, each node shows the same netgroup file hash value:

```
vserver services name-service netgroup status
```

If this is the case, you can skip the next step and proceed with the upgrade or revert. Otherwise, proceed to the next step.

- On any one node of the cluster, manually load the netgroup file:

```
vserver services netgroup load -vserver vserver_name -source uri
```

This command downloads the netgroup file on all nodes. If a netgroup file already exists on a node, it is overwritten.

Related information

[NFS management](#)

Enabling and reverting LIFs to home ports (preparing the ONTAP software for the update)

During a reboot, some LIFs might have been migrated to their assigned failover ports. Before and after you upgrade, revert, or downgrade a cluster, you must enable and revert any LIFs that are not on their home ports.

About this task

The `network interface revert` command reverts a LIF that is not currently on its home port back to its home port, provided that the home port is operational. A LIF's home port is specified when the LIF is created; you can determine the home port for a LIF by using the `network interface show` command.

Steps

- Display the status of all LIFs:

```
network interface show
```

Example

This example displays the status of all LIFs for a storage virtual machine (SVM).

```
cluster1::> network interface show -vserver vs0
-----
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
vs0	data001	down/down	192.0.2.120/24	node0	e0e	true
	data002	down/down	192.0.2.121/24	node0	e0f	true
	data003	down/down	192.0.2.122/24	node0	e2a	true
	data004	down/down	192.0.2.123/24	node0	e2b	true
	data005	down/down	192.0.2.124/24	node0	e0e	false
	data006	down/down	192.0.2.125/24	node0	e0f	false
	data007	down/down	192.0.2.126/24	node0	e2a	false
	data008	down/down	192.0.2.127/24	node0	e2b	false

```
-----
8 entries were displayed.
```

If any LIFs appear with a `Status Admin` status of down or with an `Is home` status of false, continue with the next step.

- Enable the data LIFs:

```
network interface modify {-role data} -status-admin up
```

Example

```
cluster1::> network interface modify {-role data} -status-admin up
8 entries were modified.
```

- Revert LIFs to their home ports:

```
network interface revert *
```

Example

This command reverts all LIFs back to their home ports and changes all LIF home statuses to true.

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

4. Verify that all LIFs are in their home ports:

```
network interface show
```

Example

This example shows that all LIFs for SVM vs0 are on their home ports.

```
cluster1::> network interface show -vserver vs0
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
vs0	data001	up/up	192.0.2.120/24	node0	e0e	true
	data002	up/up	192.0.2.121/24	node0	e0f	true
	data003	up/up	192.0.2.122/24	node0	e2a	true
	data004	up/up	192.0.2.123/24	node0	e2b	true
	data005	up/up	192.0.2.124/24	node1	e0e	true
	data006	up/up	192.0.2.125/24	node1	e0f	true
	data007	up/up	192.0.2.126/24	node1	e2a	true
	data008	up/up	192.0.2.127/24	node1	e2b	true

```
8 entries were displayed.
```

Preparing all load-sharing mirrors for a major upgrade

Before performing a major upgrade from ONTAP 8.3, you should move all of the load-sharing mirror source volumes to an aggregate on the node that you will upgrade last. This ensures that load-sharing mirror destination volumes are the same or later versions of ONTAP.

Steps

1. Record the locations of all load-sharing mirror source volumes.
Knowing where the load-sharing mirror source volumes came from helps facilitate returning them to their original locations after the major upgrade.
2. Determine the node and aggregate to which you will move the load-sharing mirror source volumes.
3. Move the load-sharing mirror source volumes to the node and aggregate by using the `volume move start` command.

Identifying active CIFS sessions that should be terminated

Before performing a nondisruptive upgrade or downgrade, you should identify and gracefully terminate any CIFS sessions that are not continuously available.

About this task

Continuously available CIFS shares, which are accessed by Hyper-V or Microsoft SQL Server clients using the SMB 3.0 protocol, do not need to be terminated before upgrading or downgrading.

Steps

1. Identify any established CIFS sessions that are not continuously available:
`vserver cifs session show -continuously-available !Yes -instance`

This command displays detailed information about any CIFS sessions that have no continuous availability. You should terminate them before proceeding with the ONTAP upgrade or downgrade.

Example

```
cluster1::> vserver cifs session show -continuously-available !Yes -instance

                Node: nodel
                Vserver: vs1
                Session ID: 1
                Connection ID: 4160072788
Incoming Data LIF IP Address: 198.51.100.5
Workstation IP address: 203.0.113.20
Authentication Mechanism: NTLMv2
                Windows User: CIFS\user1
                UNIX User: nobody
                Open Shares: 1
                Open Files: 2
                Open Other: 0
                Connected Time: 8m 39s
                Idle Time: 7m 45s
                Protocol Version: SMB2_1
                Continuously Available: No
1 entry was displayed.
```

2. If necessary, identify the files that are open for each CIFS session that you identified:

```
vserver cifs session file show -session-id session_ID
```

Example

```
cluster1::> vserver cifs session file show -session-id 1

Node:          nodel
Vserver:       vs1
Connection:    4160072788
Session:       1
File          File      Open Hosting      Continuously
ID           Type      Mode Volume         Share          Available
-----
1           Regular  rw  vol10             homedirshare   No
Path:       \TestDocument.docx
2           Regular  rw  vol10             homedirshare   No
Path:       \file1.txt
2 entries were displayed.
```

Related concepts

[Considerations for session-oriented protocols](#) on page 110

Configuring LDAP clients to use TLS for highest security

Before upgrading to the target ONTAP release, you must configure LDAP clients using SSLv3 for secure communications with LDAP servers to use TLS. SSL will not be available after the upgrade.

About this task

By default, LDAP communications between client and server applications are not encrypted. You must disallow the use of SSL and enforce the use of TLS.

Steps

1. Verify that the LDAP servers in your environment support TLS.

If they do not, do not proceed. You should upgrade your LDAP servers to a version that supports TLS.

2. Check which ONTAP LDAP client configurations have LDAP over SSL/TLS enabled:

```
vserver services name-service ldap client -show
```

If there are none, you can skip the remaining steps. However, you should consider using LDAP over TLS for better security.

3. For each LDAP client configuration, disallow SSL to enforce the use of TLS:

```
vserver services name-service ldap client modify -vserver vserver_name -
client-config ldap_client_config_name -allow-ssl false
```

4. Verify that the use of SSL is no longer allowed for any LDAP clients:

```
vserver services name-service ldap client show
```

Related information

[NFS management](#)

Checking for back-end configuration errors before downgrading

Before downgrading a storage system that uses array LUNs to an earlier release of ONTAP, you need to run the `storage errors show` command to determine whether there are any back-end configuration errors.

Steps

1. Check for errors that would prevent ONTAP and the back-end storage array from operating together properly:

```
storage array config show
```

- If the output *does not* instruct you to run the `storage errors show` command, there are no errors and you can proceed with the downgrade.
- If the output *does* instruct you to run the `storage errors show` command, continue with this procedure.

2. Obtain details about the error at the array LUN level:

```
storage errors show
```

Example

```
cluster1::> storage errors show
DGC_RAID5_1
-----
NAME (Serial #): This Array LUN is only available on one path.
Proper configuration requires two paths.
```

3. Fix the problems indicated by the `storage errors show` command, and then downgrade your system.

The *FlexArray virtualization installation requirements and reference guide* contains explanations about errors shown in the `storage errors show` output and provides information about how to fix them.

Related information

[FlexArray virtualization installation requirements and reference](#)

Preparing SnapMirror relationships for a nondisruptive upgrade or downgrade

You must suspend SnapMirror operations before performing a nondisruptive upgrade or downgrade of ONTAP.

Steps

1. Use the `snapmirror show` command to determine the destination path for each SnapMirror relationship.
2. For each destination volume, suspend future SnapMirror transfers:

```
snapmirror quiesce -destination-path destination
```

If there are no active transfers for the SnapMirror relationship, this command sets its status to `Quiesced`. If the relationship has active transfers, the status is set to `Quiescing` until the transfer is completed, and then the status becomes `Quiesced`.

Example

This example quiesces transfers involving the destination volume `vol1` from SVM `vs0.example.com`:

```
cluster1::> snapmirror quiesce -destination-path vs0.example.com:vol1
```

3. Verify that all SnapMirror relationships are quiesced:

```
snapmirror show -status !Quiesced
```

This command displays any SnapMirror relationships that are *not* quiesced.

Example

This example shows that all SnapMirror relationships are quiesced:

```
cluster1::> snapmirror show -status !Quiesced
There are no entries matching your query.
```

4. If any SnapMirror relationships are currently being transferred, do one of the following options:

Option	Description
Wait for the transfers to finish before performing the ONTAP upgrade.	After each transfer finishes, the relationship changes to <code>Quiesced</code> status.
Stop the transfers\: snapmirror abort - destination-path destination -h	This command stops the SnapMirror transfer and restores the destination volume to the last Snapshot copy that was successfully transferred. The relationship is set to <code>Quiesced</code> status.
Note: You must use the <code>-foreground true</code> parameter if you are aborting load-sharing mirror transfers.	

Related concepts

[Upgrade requirements for SnapMirror](#) on page 12

Preparing to upgrade nodes using NetApp Storage Encryption with external key management servers

If you are using NetApp Storage Encryption (NSE) and upgrading to ONTAP 9.3 or later, you must delete any existing external key management (KMIP) server connections before performing the upgrade.

Steps

1. Verify that the NSE drives are unlocked, open, and set to the default manufacture secure ID 0x0:
`storage encryption disk show -disk*`
2. Enter the advanced privilege mode:
`set -privilege advanced`
3. Use the default manufacture secure ID 0x0 to assign the FIPS key to the self-encrypting disks (SEDs):
`storage encryption disk modify -fips-key-id 0x0 -disk *`
4. Verify that assigning the FIPS key to all disks is complete:
`storage encryption disk show-status`
5. Verify that the **mode** for all disks is set to data:
`storage encryption disk show`
6. View the configured KMIP servers:
`security key-manager show`
7. Delete the configured KMIP servers:
`security key-manager delete -address kmip_ip_address`
8. Delete the external key manager configuration:
`security key-manager delete-kmip-config`
Note: This step does not remove the NSE certificates

After you finish

After the upgrade is complete, you must reconfigure the KMIP server connections.

Related tasks

[Reconfiguring KMIP server connections after upgrading to ONTAP 9.3 or later](#) on page 78

After performing an upgrade to ONTAP 9.3 or later, you must reconfigure your external key management (KMIP) server connections.

Obtaining ONTAP software images

For ONTAP 9.4 and later, you can copy the ONTAP software image from the NetApp Support Site to a local folder. For upgrades from ONTAP 9.3 or earlier, you must copy the ONTAP software image to an HTTP server or FTP server on your network.

About this task

To upgrade, revert, or downgrade the cluster to the target release of ONTAP, you require access to software images. Software images, firmware version information, and the latest firmware for your

platform model are available on the NetApp Support Site. You should note the following important information:

- Software images are specific to platform models.
You must obtain the correct image for your cluster.
- Software images include the latest version of system firmware that was available when a given version of ONTAP was released.
- If you are upgrading a system with NetApp Volume Encryption to ONTAP 9.5 or later, you must download the ONTAP software image for non-restricted countries, which includes NetApp Volume Encryption.
If you use the ONTAP software image for restricted countries to upgrade a system with NetApp Volume Encryption, the system panics and you lose access to your volumes.

Steps

1. Locate the target ONTAP software in the **Software Downloads** area of the NetApp Support Site.
2. Copy the software image.
 - For ONTAP 9.3 or earlier, copy the software image (for example, `93_q_image.tgz`) from the NetApp Support Site to the directory on the HTTP server or FTP server from which the image will be served
 - For ONTAP 9.4 or later, copy the software image (for example, `95_q_image.tgz`) from the NetApp Support Site to the directory on the HTTP server or FTP server from which the image will be served or to a local folder.

Related information

[NetApp Downloads: Software](#)

Installing the ONTAP software image

You must install the target software image on the cluster's nodes.

Before you begin

Note: If you are upgrading a system with NetApp Volume Encryption to ONTAP 9.5 or later, you must have downloaded the ONTAP software image for non-restricted countries, which includes NetApp Volume Encryption.

If you use the ONTAP software image for restricted countries to upgrade a system with NetApp Volume Encryption, the system panics and you lose access to your volumes.

Steps

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


```
set -privilege advanced
```

The advanced prompt (`*>`) appears.
2. Install the software image on the nodes:


```
system node image update -node * -package location -replace-package true -setdefault true -background true
```

This command downloads and installs the software image on all of the nodes simultaneously. To download and install the image on each node one at a time, do not specify the `-background` parameter.

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

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

This command displays the current status of the software image download and 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.

Example

This example shows a two-node cluster in which the software image is downloaded and 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.
```

Selecting your update method for non-MetroCluster configurations

Based the requirements of your non-MetroCluster configuration, you can update (upgrade or downgrade) a cluster to a different ONTAP release by performing a nondisruptive upgrade or a disruptive upgrade.

Note: This topic describes options for updating non-MetroCluster configurations. If you are upgrading a MetroCluster configuration, see [Selecting your update method for MetroCluster configurations](#) on page 35.

Nondisruptive and disruptive updates

Nondisruptive upgrade and downgrade procedures perform the operation while maintaining service to clients.

In a *disruptive upgrade or downgrade*, storage failover is disabled for each HA pair, and then each node is rebooted one at a time. Disruptive upgrades can be performed more quickly than nondisruptive upgrades, and require fewer steps to complete. 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 upgrade or downgrade is always used for single-node clusters.

Automated and manual updates

Automated nondisruptive upgrades (NDU) are the preferred method of upgrading a cluster. With NDU, 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 the upgrade in the background.

Nondisruptive manual upgrades involve manual steps to confirm the ONTAP configuration on each node and then use the rolling update method to perform the upgrade or downgrade. In the rolling update method, a node is taken offline and updated 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, this time on the partner node. Each additional HA pair is upgraded in sequence until all HA pairs are running the target release

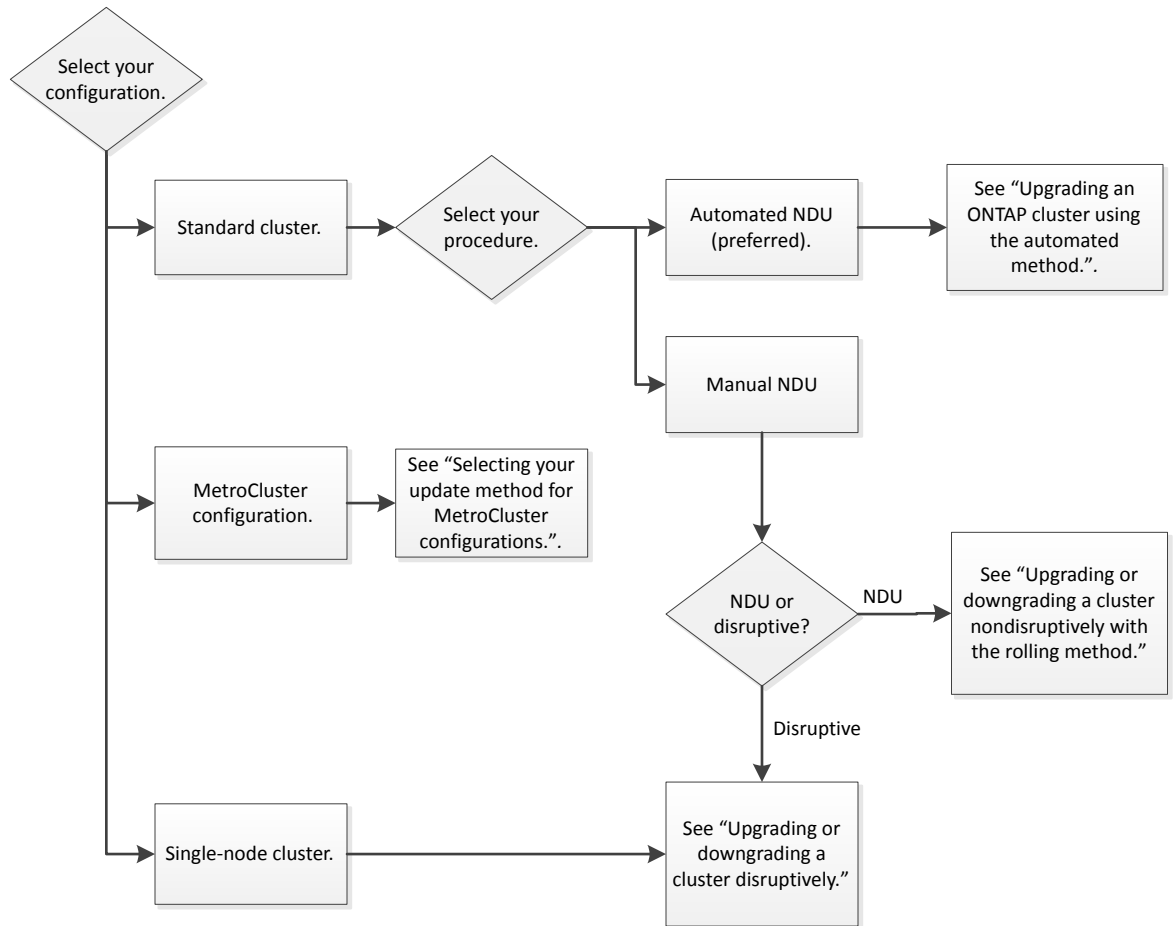
Starting with ONTAP 9.2, automatic updates can also be performed on single-node clusters. However, because single-nodes lack redundancy, updates are disruptive.

Choosing the right procedure

You can use the following diagram to determine which procedure you should use, based on the following criteria:

- Your cluster configuration
- Whether you choose nondisruptive or disruptive procedures

Below the diagram are links to the procedures.



For this upgrade or downgrade procedure	Nondisruptive?	Manual?	See...
Automated	No for single-node clusters; yes for all others.	No	Upgrading an ONTAP cluster using the automated method on page 36
Rolling	Yes	Yes	Upgrading or downgrading a cluster nondisruptively by using the rolling upgrade method on page 40
Disruptive	No	Yes	Updating an ONTAP cluster disruptively on page 69

Selecting your update method for MetroCluster configurations

Based on your requirements, you can update (upgrade or downgrade) a MetroCluster configuration to a different ONTAP release by performing a nondisruptive upgrade or a disruptive upgrade.

Nondisruptive and disruptive updates

Nondisruptive upgrade and downgrade procedures perform the operation while maintaining service to clients.

In a *disruptive upgrade or downgrade*, storage failover is disabled for each HA pair, and then each node is rebooted one at a time. Disruptive upgrades can be performed more quickly than nondisruptive upgrades, and require fewer steps to complete. 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.

Automated and manual updates

Automated nondisruptive upgrades (NDU) are the preferred method of upgrading a cluster. With NDU, 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 the upgrade in the background.

Choosing the right procedure

You can use the following table to determine which procedure you should use and whether that procedure is manual or automated, based on the following criteria:

- The number of nodes in the MetroCluster configuration
- The ONTAP version you are using
- Whether you choose nondisruptive or disruptive procedures

Number of nodes in MetroCluster configuration	ONTAP version	Nondisruptive?	Manual?	See...
Two	9.3 and later	Yes	No	Upgrading an ONTAP cluster using the automated method on page 36
Two	9.2 and earlier	No for downgrades from ONTAP 9.1; yes for all others	Yes	Updating a two-node MetroCluster configuration in ONTAP 9.2 or earlier on page 65 Downgrading a two-node MetroCluster configuration disruptively on page 68
Four	9.3 and later	Yes	No	Upgrading an ONTAP cluster using the automated method on page 36
Four or eight	9.2 and earlier	Yes	Yes	Updating a four- or eight-node MetroCluster configuration manually on page 51
Four or eight, patch upgrades only	Any	Yes	No	Software express upgrade
Eight	9.3 and later	Yes	Yes	Updating a four- or eight-node MetroCluster configuration manually on page 51
Any	Any	No	Yes	Updating an ONTAP cluster disruptively on page 69

Upgrading an ONTAP cluster using the automated method

The automated upgrade method validates the cluster components to verify that the cluster can be upgraded, installs the target ONTAP image on each node, and then executes the upgrade in the background. Automated upgrades of multi-node clusters are non-disruptive. Automated upgrades of single-node clusters are disruptive because single-node clusters lack redundancy.

Requesting notification of issues encountered in nondisruptive upgrades

If you do not plan to monitor the progress of the upgrade process, it is a good practice to request EMS notifications of errors that might require manual intervention. Alternatively, you can configure an AutoSupport message to send to your internal support organization.

Before you begin

You must be a cluster administrator to perform this task.

About this task

It is useful to set up notifications such that they are sent in case of a problem during the upgrade process. In particular, the `callhome.andu.pausederr` message contains useful troubleshooting information.

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

Steps

1. Request notification of issues encountered in nondisruptive upgrade.

EMS express configuration

2. Before initiating a nondisruptive upgrade, if AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message. If AutoSupport is not enabled on this cluster, then ignore this step: **system node autosupport invoke -node * -type all -message MAINT=xh**

x is the duration of the maintenance window in hours.

Note: The message will notify technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

Example

This command suppresses automatic case creation for two hours:

```
cluster::*> system node autosupport invoke -node * -type all -message
MAINT=2h
```

3. After the nondisruptive upgrade, reenables automatic case creation by invoking an AutoSupport message: **system node autosupport invoke -node * -type all -message MAINT=END**

Example

The command reenables automatic case creation:

```
cluster::*> system node autosupport invoke -node * -type all -message
MAINT=END
```

Related information

[ONTAP 9 commands](#)

[EMS express configuration](#)

Performing an automatic nondisruptive upgrade using the CLI

You can use the command line interface (CLI) to verify that the cluster can be upgraded nondisruptively, install the target ONTAP image on each node, and then, execute an upgrade in the background.

Before you begin

- You must have met the upgrade preparation requirements.
- For each HA pair, each node should have one or more ports on the same broadcast domain. When a set of nodes is upgraded during a batch upgrade, the LIFs are migrated to the HA partner nodes. If the partners do not have any ports in the same broadcast domain, then the LIF migration fails.

About this task

The `cluster image validate` command checks the cluster components to validate that the upgrade can be completed nondisruptively, and then provides the status of each check and any required action you must take before performing the software upgrade.

Steps

1. Delete the previous ONTAP software package:

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

2. Download the target ONTAP software package:

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

Example

```
cluster1::> cluster image package get -url http://www.example.com/software/9.6/
image.tgz

Package download completed.
Package processing completed.
```

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

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

Example

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

4. Verify that the cluster is ready to be upgraded nondisruptively:

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

If you are upgrading a two-node or four-node MetroCluster configuration, you must run this command on all nodes before proceeding.

Example

```
cluster1::> cluster image validate -version 9.6

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

5. Monitor the progress of the validation:

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

6. Complete all required actions identified by the validation.

7. Generate a software upgrade estimate:

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

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

8. Perform the software upgrade:

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

If the cluster consists of 2 through 6 nodes, a rolling upgrade is performed.

If the cluster consists of 8 or more nodes, a batch upgrade is performed by default. If desired, you can use the `-force-rolling` parameter to specify a rolling upgrade instead.

After completing each takeover and each 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.

Example

```
cluster1::> cluster image update -version 9.6

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::>
```

9. Display the cluster update progress:

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

Note: 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.

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

Example

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

Update Phase	Status	Estimated Duration	Elapsed 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.

```
cluster1::>
```

11. 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

Resuming an upgrade after an error in the automated upgrade process

If an automated upgrade pauses because of an error, you can resolve the error and resume the automated upgrade, or you can cancel the automated upgrade and complete the process manually. If you choose to continue the automated upgrade, do not perform any of the upgrade steps manually.

Steps

1. Depending on the System Manager version that you are running, perform one of the following steps:
 - ONTAP 9.4 or earlier: Click **Configuration > Cluster Update**.
 - ONTAP 9.5 or later: Click **Configuration > Cluster > Update**.
2. Continue the automated update or cancel it and continue manually.

If you want to...	Then...
Resume the automated updated	Click Resume .
Cancel the automated updated and continue manually	Click Cancel .

Upgrading or downgrading a cluster nondisruptively by using the rolling upgrade method

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.

Before you begin

You must have satisfied upgrade or downgrade preparation requirements.

About this task

The versions used in these task examples might vary depending on whether you are upgrading or downgrading the software version, or if you are performing a major or minor upgrade or downgrade.

Steps

1. [Upgrading the first node in an HA pair](#) on page 41
You upgrade or downgrade 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. [Upgrading the second node in an HA pair](#) on page 45
After upgrading the first node in an HA pair, you upgrade or downgrade its partner by initiating a takeover on it. The first node serves the partner's data while the partner node is upgraded.
3. Repeat [1](#) on page 41 and [2](#) on page 41 for each additional HA pair.

After you finish

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 or downgraded.

About this task

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.

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:

```
cluster image show-update-progress
```
5. Verify that the new ONTAP software image is set as the default image:

```
system image show
```

Example

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      Version  Install
-----  -
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
```

Example

```
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.

Example

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.

Example

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
vservers 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.

Example

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. 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.

- b. 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.

Example

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

```
cluster1::> network port show -node node0
Node  Port      IPspace      Broadcast Domain Link  MTU      Speed (Mbps)
-----  -
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.

Example

```
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 -data-protocol nfs|cifs -role data -curr-node
nodenameA
```

Example

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 -data-protocol nfs|cifs -role data -curr-node node0
-----
Vserver      Logical      Status      Network      Current      Current Is
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.

Example

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
```

After you finish

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 or downgraded.

Steps

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:

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

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

```
system image show
```

Example

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      Is      Install
          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   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
```

Example

```
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.

Example

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.

Example

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
vservers 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
```

Example

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:

- 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:

```
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.

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.

Example

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

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

Node	Port	IPspace	Broadcast	Domain	Link	MTU	Speed (Mbps) 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.

Example

```

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 -data-protocol nfs|cifs -role data -curr-node nodenameB
```

Example

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 -data-protocol nfs|cifs -role data -curr-node node1
Vserver      Logical      Status      Network      Current      Current      Is
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.

Example

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
```

Example

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      Is      Version  Install
-----  -----  -----  -----  -----  -----
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`

After you finish

Upgrade any additional HA pairs.

Updating a MetroCluster configuration using the manual method

For the nondisruptive upgrade or downgrade of some MetroCluster configurations, you must use the manual procedure. The procedure used depends on the number of nodes in the MetroCluster configuration and the ONTAP version. The procedures apply to both MetroCluster FC and MetroCluster IP configurations.

About this task

To determine the upgrade procedure you should use, see [Selecting your update method for MetroCluster configurations](#) on page 35.

Choices

- [Downgrade requirements for MetroCluster configurations](#) on page 50
- [Updating a four- or eight-node MetroCluster configuration manually](#) on page 51
- [Updating a two-node MetroCluster configuration in ONTAP 9.2 or earlier](#) on page 65
- [Downgrading a two-node MetroCluster configuration disruptively](#) on page 68

Downgrade requirements for MetroCluster configurations

You should be aware of some important requirements when downgrading MetroCluster configurations.

General requirements

- Both clusters must be running the same version of ONTAP.
You can verify the ONTAP version by using the `version` command.

- Eight-node MetroCluster configurations can be downgraded non-disruptively from ONTAP 9.1 to 9.0.
An eight-node MetroCluster configuration cannot be reverted to the Data ONTAP 8.3.x.
- A two-node MetroCluster configuration can only be disruptively downgraded from ONTAP 9.1 to 9.0.
Downgrading a two-node MetroCluster configuration disruptively on page 68
- Eight-node or four-node MetroCluster configurations must be downgraded using the lockstep procedure in which DR pairs are downgraded simultaneously.
Updating a four- or eight-node MetroCluster configuration manually on page 51
- The MetroCluster configuration must be in either normal mode or switchover mode.
- The aggregates in both clusters must not be in `resyncing` RAID status.
During MetroCluster healing, the mirrored aggregates are resynchronized. You can verify whether the MetroCluster configuration is in this state by using the `storage aggregate plex show -in-progress true` command. If any aggregates are in progress, the resynchronization process is still underway and you should not perform a downgrade until the aggregate resynchronization is complete.
- Negotiated switchover operations fail while the downgrade is in progress.
After the downgrade has started, you should not attempt a negotiated switchover until both clusters have been downgraded, and all nodes are running the same version of ONTAP. If a site failure occurs during the downgrade, you should perform a forced switchover.
- The MetroCluster operation history might not be available after the downgrade.
If you previously used the `metrocluster check run` command while running the higher version of ONTAP, then after the downgrade, the `metrocluster operation show` and `metrocluster operation history show` commands incorrectly display “12” instead of the previous check operation.

Configuration requirements for MetroCluster configurations in normal operation

- The source storage virtual machine (SVM) LIFs must be up and located on their home nodes. Data LIFs for the destination SVMs are not required to be up or to be on their home nodes.
- All aggregates at the local site must be online.
- All root and data volumes that are owned by the SVMs of the local cluster must be online.

Configuration requirements for MetroCluster configurations in switchover

- All LIFs must be up and located on their home nodes.
- All aggregates must be online, except for the root aggregates at the disaster recovery (DR) site. Root aggregates at the DR site are offline during certain phases of a switchover.
- All volumes must be online.


Updating a four- or eight-node MetroCluster configuration manually

The manual update procedure for upgrading or downgrading 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 some post-update tasks.

About this task

- 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 software on an eight-node or four-node MetroCluster configuration

The MetroCluster software update 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:



The MetroCluster software update procedure involves upgrading or downgrading one DR group at a time.

For four-node MetroCluster configurations:

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

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

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

Preparing to update a MetroCluster DR group

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

Before you begin

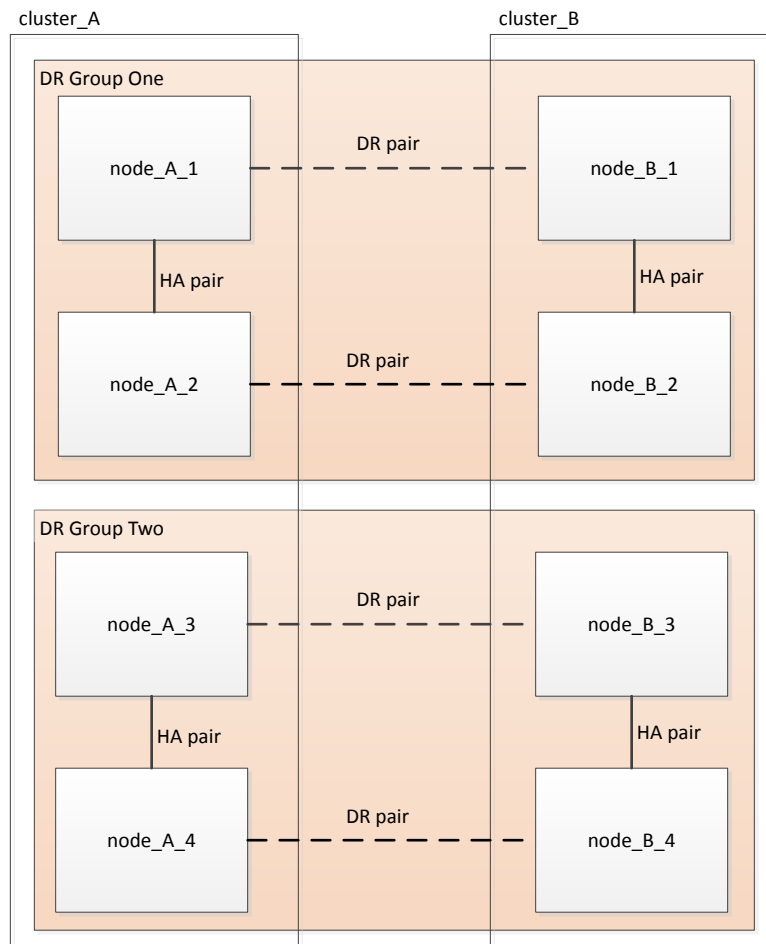
You must have installed the software images.

[Installing the ONTAP software image](#) on page 32

About this task

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:



Steps

1. Identify the DR pairs in the configuration:

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

Example

```
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 running on each node:

- a. Confirm the version on cluster_A:

```
system image show
```

Example

```
cluster_A::*> system image show
Node      Image      Is      Is      Version      Install
           Image    Default Current  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::*>
```

- b. Confirm the version on cluster_B:

```
system image show
```

Example

```
cluster_B::*> system image show
Node      Image      Is      Is      Version      Install
           Image    Default Current  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::>

```

4. Trigger an AutoSupport notification:

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

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

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

5. 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.

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

- a. Verify the images on cluster_A:

```
system image show
```

Example

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
      Is      Is      Install
Node  Image  Default Current Version 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.

```

- b. Verify the images on cluster_B:

```
system image show
```

Example

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
      Is      Is      Install
Node  Image  Default Current Version 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.

```

7. Determine whether the nodes to be upgraded are currently serving any clients by entering the following command 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.

Note: 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.

Example

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.

Before you begin

All nodes must be running the old version of ONTAP.

About this task

In this task, `node_A_1` and `node_B_1` are updated.

If you have updated 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 would be updating `node_A_3` and `node_B_3`.

Steps

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
```

Example

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. Ensure 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.

Note: 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
```

Example

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.

Note: 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
```

Example

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:

- 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. 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
```

Example

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

```

cluster_A::*> system image show
      Is      Is      Install
Node  Image  Default Current Version  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
```

Example

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
      Is      Is      Install
Node  Image  Default Current Version  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.

Before you begin

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

About this task

In this task, node_A_2 and node_B_2 are updated.

If you have updated 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.

Steps

1. 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:

If you are upgrading from ...	Enter this command...
ONTAP 9.1	<code>storage failover takeover -ofnode node_A_2</code>
ONTAP 9.0 or Data ONTAP 8.3.x	<code>storage failover takeover -ofnode node_A_2 -option allow-version-mismatch</code> 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.

The node boots up to the `waiting for giveback` state.

Note: 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
```

Example

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.

Example

```
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.
```

2. 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	<code>storage failover takeover -ofnode node_B_2</code>
ONTAP 9.0 or Data ONTAP 8.3.x	<code>storage failover takeover -ofnode node_B_2 -option allow-version-mismatch</code> 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.

The node boots up to the `waiting for giveback` state.

Note: 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**Example**

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.
```

3. 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.

4. 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.

5. 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.

6. 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.

[High-availability configuration](#)

- 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`.

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 giveback.

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

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

```
set -privilege advanced
```

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

9. Confirm the version on cluster_A:

```
system image show
```

Example

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          Is          Version    Install
          Image    Default    Current    Version    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::>
```

10. Confirm the version on cluster_B:

```
system image show
```

Example

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          Is          Version    Install
          Image    Default    Current    Version    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::>
```

11. 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.

12. Verify that automatic giveback is enabled:

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


Example

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.
```

Updating a two-node MetroCluster configuration in ONTAP 9.2 or earlier

You can upgrade and in some cases downgrade ONTAP nondisruptively for a two-node MetroCluster configuration. This method has several steps: 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.

About this task

- This procedure is for two-node MetroCluster configurations running ONTAP 9.2 or earlier only. Do not use this procedure if you have a four-node MetroCluster configuration.
- For downgrades, this procedure is only for downgrading from ONTAP 9.0 or earlier. You cannot use this procedure to downgrade a two-node MetroCluster configuration from ONTAP 9.1 or ONTAP 9.2, which can only be done disruptively.

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
```

Example

```
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
```

Example

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

```
cluster_B::*> system node image show
Node      Image      Is      Is      Install
          Image      Default Current Version      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.

Example

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 vserver 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.

Example

```
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
```

Example

```
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
```

Example

```
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
```

Example

```
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 vserver show
```

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

```
metrocluster switchback
```

15. Verify that the switchback was completed successfully:

```
metrocluster operation show
```

Example

```
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
```

Example

```
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
```

Example

```
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

You should perform any post-upgrade or post-downgrade tasks.

Related information

[MetroCluster management and disaster recovery](#)

Downgrading a two-node MetroCluster configuration disruptively

A two-node MetroCluster (MCC) configuration can only be downgraded from ONTAP 9.1 to ONTAP 9 disruptively.

About this task

A two-node MCC configuration cannot be downgraded from ONTAP 9.2 to ONTAP 9.1. You can only revert from ONTAP 9.2 to an earlier version of ONTAP.

Steps

1. Disable automatic unplanned switchover (AUSO) on both the clusters:

```
metrocluster modify -auto-switchover-failure-domain auso-disabled
```

2. Verify that AUSO is disabled:

```
metrocluster show
```

AUSO Failure Domain is auso-disabled.

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

```
set -privilege advanced
```

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

4. Perform the following steps on site A, and then repeat the same steps on site B.

- a. Install the ONTAP 9 software image and set it as the default:

```
system node image update -package package_location
```

- b. Disable the new features and capabilities that are not available in ONTAP 9.

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

```
system node image show
```

Example

The following example shows that the 9.0 image is set as the default image:

```
cluster_B::*> system node image show
      Is      Is
Node  Image  Default Current Version
-----
node_B_1
      image1  false   true   9.1
      image2  true    false  9.0
2 entries were displayed.
```

- d. Reboot the node:

```
system node reboot -node nodename
```

- e. After the reboot is complete, verify that the storage virtual machines (SVMs) are running and the LIFS are online:

```
network interface show -vserver vservice_name
```

5. Enable AUSO on both clusters:

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

6. Verify that AUSO is enabled:

```
metrocluster show
```

AUSO Failure Domain is auso-on-cluster-disaster.

7. Validate the configuration:

```
metrocluster check
```

Updating an ONTAP cluster disruptively

If you can take your cluster offline to upgrade or downgrade 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.

Before you begin

- You must have satisfied preparation requirements.
- If you are operating in a SAN environment, all SAN clients must be shut down or suspended until the upgrade or downgrade is complete.

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

About this task

In a disruptive upgrade or downgrade, 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
```

Example

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      Install
-----  -----
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.
Two nodes	<ol style="list-style-type: none"> a. Disable cluster high availability: <code>cluster ha modify -configured false</code> b. Disable storage failover for the HA pair: <code>storage failover modify -node * -enabled false</code>
More than two nodes	Disable storage failover for each HA pair in the cluster: <code>storage failover modify -node * -enabled false</code>

5. Reboot a node in the cluster:

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

Attention: 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, confirm that the new software is running:

```
system node image show
```

Example

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          Is          Install
-----  -
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.
```

7. Verify that the upgrade or downgrade is completed successfully:

- a. Set the privilege level to advanced:

```
set -privilege advanced
```

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

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

The status should be listed as `complete`.

If the upgrade or downgrade is not successful, from the node, run the `system node upgrade-revert upgrade` command. If this command does not complete the node's upgrade or downgrade, contact technical support immediately.

- c. Return to the admin privilege level:

```
set -privilege admin
```

8. Repeat Steps 5 on page 70 through 7 on page 71 for each additional node.
9. 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
```

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

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

Performing an automated upgrade on a single-node cluster

Beginning with ONTAP 9.2, you can perform an automated update of a single-node cluster. Because single-node clusters lack redundancy, updates are always disruptive.

Before you begin

- You must have satisfied upgrade preparation requirements.

Steps

1. Delete the previous ONTAP software package:

```
cluster image package delete -version previous_package_version
```

2. Download the target ONTAP software package:

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

Example

```
cluster1::> cluster image package get -url http://www.example.com/software/9.6/
image.tgz

Package download completed.
Package processing completed.
```

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

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

Example

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

4. Verify that the cluster is ready to be upgraded:

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

Example

```
cluster1::> cluster image validate -version 9.6

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

5. Monitor the progress of the validation:

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

6. Complete all required actions identified by the validation.

7. Optionally, generate a software upgrade estimate:

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

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

8. Perform the software upgrade:

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

Note: If an issue is encountered, the update pauses and prompts you to take corrective action. You can use the `cluster image show-update-progress` command to view details about any issues and the progress of the update. After correcting the issue, you can resume the update by using the `cluster image resume-update` command.

9. Display the cluster update progress:

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

The node is rebooted as part of the update and cannot be accessed while rebooting.

10. Trigger a notification:

```
autosupport invoke -node * -type all -message "Finishing Upgrade"
```

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

Completing post-upgrade or downgrade tasks for the cluster

After you upgrade or downgrade a cluster to a different version of ONTAP software, you must complete additional tasks to restore normal operation.

Steps

1. [Verifying the cluster version](#) on page 74
After all of the HA pairs have been upgraded, you must use the `version` command to verify that all of the nodes are running the target release.
2. [Verifying cluster health \(verifying storage health\)](#) on page 74
Before and after you upgrade, revert, or downgrade a cluster, you should verify that the nodes are healthy and eligible to participate in the cluster, and that the cluster is in quorum.
3. [Verifying storage health \(completing post-upgrade or downgrade tasks\)](#) on page 75
Before and after you upgrade, revert, or downgrade a cluster, you should verify the status of your disks, aggregates, and volumes.
4. [Verifying networking and storage status for MetroCluster configurations \(post-upgrade or downgrade\)](#) on page 76
Before and after performing an update in a MetroCluster configuration, you should verify the status of the LIFs, aggregates, and volumes for each cluster.
5. [Verifying the SAN configuration after an upgrade](#) on page 78
If you are upgrading in a SAN environment, then after the upgrade, you should verify that each initiator that was connected to a LIF before the upgrade has successfully reconnected to the LIF.
6. [Reconfiguring KMIP server connections after upgrading to ONTAP 9.3 or later](#) on page 78
After performing an upgrade to ONTAP 9.3 or later, you must reconfigure your external key management (KMIP) server connections.
7. [Enabling and reverting LIFs to home ports \(post-upgrade or downgrade tasks for the cluster\)](#) on page 79
During a reboot, some LIFs might have been migrated to their assigned failover ports. Before and after you upgrade, revert, or downgrade a cluster, you must enable and revert any LIFs that are not on their home ports.
8. [Relocating moved load-sharing mirror source volumes](#) on page 80
After successfully completing a nondisruptive upgrade, you can move load-sharing mirror source volumes back to the locations they were in originally before the upgrade.
9. [Resuming SnapMirror operations](#) on page 80
After completing a nondisruptive upgrade or downgrade, you must resume any SnapMirror relationships that were suspended.
10. [Setting the desired NT ACL permissions display level for NFS clients](#) on page 81
After upgrading from ONTAP 8.3.0, the default handling for displaying NT ACL permissions to NFS clients has changed. You should check the setting and change it to the desired setting for your environment if necessary. This task does not apply if you are upgrading from ONTAP 8.3.1 or later.
11. [Enforcing SHA-2 on administrator account passwords](#) on page 82
Administrator accounts created prior to ONTAP 9.0 continue to use MD5 passwords after the upgrade, until the passwords are manually changed. MD5 is less secure than SHA-2. Therefore, after upgrading, you should prompt users of MD5 accounts to change their passwords to use the default SHA-512 hash function.
12. [When you need to update the Disk Qualification Package](#) on page 83

The Disk Qualification Package (DQP) adds full support for newly qualified drives. Before you update drive firmware or add new drive types or sizes to a cluster, you must update the DQP. A best practice is to also update the DQP regularly; for example, every quarter or semi-annually.

Verifying the cluster version

After all of the HA pairs have been upgraded, you must use the `version` command to verify that all of the nodes are running the target release.

About this task

The cluster version is the lowest version of ONTAP running on any node in the cluster. If the cluster version is not the target ONTAP release, you can upgrade your cluster.

Steps

1. Verify that the cluster version is the target ONTAP release:
`version`
2. If the cluster version is not the target ONTAP release, you can verify the upgrade status of all nodes
`system node upgrade-revert show`

Verifying cluster health (verifying storage health)

Before and after you upgrade, revert, or downgrade a cluster, you should verify that the nodes are healthy and eligible to participate in the cluster, and that the cluster is in quorum.

Steps

1. Verify that the nodes in the cluster are online and are eligible to participate in the cluster:
`cluster show`

Example

```
cluster1::> cluster show
Node                Health  Eligibility
-----
node0                true   true
node1                true   true
```

If any node is unhealthy or ineligible, check EMS logs for errors and take corrective action.

2. Set the privilege level to advanced:
`set -privilege advanced`
3. Enter `y` to continue.
4. Verify the configuration details for each RDB process.
 - The relational database epoch and database epochs should match for each node.
 - The per-ring quorum master should be the same for all nodes.
Note that each ring might have a different quorum master.

To display this RDB process...	Enter this command...
Management application	<code>cluster ring show -unitname mgmt</code>
Volume location database	<code>cluster ring show -unitname vldb</code>
Virtual-Interface manager	<code>cluster ring show -unitname vifmgr</code>
SAN management daemon	<code>cluster ring show -unitname bcomd</code>

Example

This example shows the volume location database process:

```
cluster1::*> cluster ring show -unitname vldb
Node      UnitName Epoch   DB Epoch DB Trnxs Master   Online
-----
node0     vldb     154     154     14847  node0   master
node1     vldb     154     154     14847  node0   secondary
node2     vldb     154     154     14847  node0   secondary
node3     vldb     154     154     14847  node0   secondary
4 entries were displayed.
```

- If you are operating in a SAN environment, verify that each node is in a SAN quorum:

```
event log show -messagename scsiblade.*
```

The most recent `scsiblade` event message for each node should indicate that the `scsi-blade` is in quorum.

Example

```
cluster1::*> event log show -messagename scsiblade.*
Time      Node      Severity  Event
-----
MM/DD/YYYY TIME  node0     INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
MM/DD/YYYY TIME  node1     INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
```

- Return to the admin privilege level:

```
set -privilege admin
```

Related information

[System administration](#)

Verifying storage health (completing post-upgrade or downgrade tasks)

Before and after you upgrade, revert, or downgrade a cluster, you should verify the status of your disks, aggregates, and volumes.

Steps

- If you are preparing to upgrade, revert, or downgrade, verify disk status:

To check for...	Do this...
Broken disks	<ol style="list-style-type: none"> Display any broken disks: <code>storage disk show -state broken</code> Remove or replace any broken disks.

To check for..	Do this..
Disks undergoing maintenance or reconstruction	<ol style="list-style-type: none"> a. Display any disks in maintenance, pending, or reconstructing states: storage disk show -state maintenance pending reconstructing b. Wait for the maintenance or reconstruction operation to finish before proceeding.

2. Verify that all aggregates are online by displaying the state of physical and logical storage, including storage aggregates:

storage aggregate show -state !online

This command displays the aggregates that are *not* online. All aggregates must be online before and after performing a major upgrade or reversion.

Example

```
cluster1::> storage aggregate show -state !online
There are no entries matching your query.
```

3. Verify that all volumes are online by displaying any volumes that are *not* online:

volume show -state !online

All volumes must be online before and after performing a major upgrade or reversion.

Example

```
cluster1::> volume show -state !online
There are no entries matching your query.
```

4. Verify that there are no inconsistent volumes:

volume show -is-inconsistent true

If any inconsistent volumes are returned, you must contact NetApp Support before you proceed with the upgrade.

Related information

[Disk and aggregate management](#)

Verifying networking and storage status for MetroCluster configurations (post-upgrade or downgrade)

Before and after performing an update in a MetroCluster configuration, you should verify the status of the LIFs, aggregates, and volumes for each cluster.

Steps

1. Verify the LIF status:

network interface show

In normal operation, LIFs for source SVMs must have an admin status of up and be located on their home nodes. LIFs for destination SVMs are not required to be up or located on their home nodes. In switchover, all LIFs have an admin status of up, but they do not need to be located on their home nodes.

Example

```
cluster1::> network interface show
Logical      Status      Network      Current      Current Is
Vserver      Interface  Admin/Oper  Address/Mask  Node         Port   Home
-----
Cluster
cluster1-a1_clus1
                up/up      192.0.2.1/24  cluster1-01
                e2a       true
cluster1-a1_clus2
                up/up      192.0.2.2/24  cluster1-01
                e2b       true
cluster1-01
clus_mgmt      up/up      198.51.100.1/24  cluster1-01
                e3a       true
cluster1-a1_inet4_intercluster1
                up/up      198.51.100.2/24  cluster1-01
                e3c       true
...
27 entries were displayed.
```

2. Verify the state of the aggregates:

```
storage aggregate show -state !online
```

This command displays any aggregates that are *not* online. In normal operation, all aggregates located at the local site must be online. However, if the MetroCluster configuration is in switchover, root aggregates at the disaster recovery site are permitted to be offline.

Example

This example shows a cluster in normal operation:

```
cluster1::> storage aggregate show -state !online
There are no entries matching your query.
```

Example

This example shows a cluster in switchover, in which the root aggregates at the disaster recovery site are offline:

```
cluster1::> storage aggregate show -state !online
Aggregate      Size Available Used% State      #Vols  Nodes      RAID Status
-----
aggr0_b1
                0B         0B      0% offline    0 cluster2-01  raid_dp,
                mirror
                degraded
aggr0_b2
                0B         0B      0% offline    0 cluster2-02  raid_dp,
                mirror
                degraded
2 entries were displayed.
```

3. Verify the state of the volumes:

```
volume show -state !online
```

This command displays any volumes that are *not* online.

If the MetroCluster configuration is in normal operation (it is not in switchover state), the output should show all volumes owned by the cluster's secondary SVMs (those with the SVM name appended with "-mc").

Those volumes come online only in the event of a switchover.

Example

This example shows a cluster in normal operation, in which the volumes at the disaster recovery site are not online.

```
cluster1::> volume show -state !online
(volume show)
Vserver   Volume      Aggregate   State   Type   Size   Available   Used%
-----
vs2-mc    vol1        aggr1_b1    -       RW     -       -           -
vs2-mc    root_vs2    aggr0_b1    -       RW     -       -           -
vs2-mc    vol2        aggr1_b1    -       RW     -       -           -
vs2-mc    vol3        aggr1_b1    -       RW     -       -           -
vs2-mc    vol4        aggr1_b1    -       RW     -       -           -
5 entries were displayed.
```

4. Verify that there are no inconsistent volumes:

```
volume show -is-inconsistent true
```

If any inconsistent volumes are returned, you must contact NetApp Support before you proceed with the upgrade.

Verifying the SAN configuration after an upgrade

If you are upgrading in a SAN environment, then after the upgrade, you should verify that each initiator that was connected to a LIF before the upgrade has successfully reconnected to the LIF.

Step

1. Verify that each initiator is connected to the correct LIF.

You should compare the list of initiators to the list you made during the upgrade preparation.

For...	Enter...
iSCSI	<code>iscsi initiator show -fields igroup,initiator-name,tpgroup</code>
FC	<code>fcp initiator show -fields igroup,wwpn,lif</code>

Reconfiguring KMIP server connections after upgrading to ONTAP 9.3 or later

After performing an upgrade to ONTAP 9.3 or later, you must reconfigure your external key management (KMIP) server connections.

Steps

1. Configure the key manager connectivity:

```
security key-manager setup
```

2. Add your KMIP servers:

```
security key-manager add -address key_management_server_ip_address
```

3. Verify that KMIP servers are connected:

```
security key-manager show -status
```

4. Query the key servers:

```
security key-manager query
```

5. Create a new authentication key and passphrase:

```
security key-manager create-key -prompt-for-key true
```

The passphrase must have a minimum of 32 characters.

6. Query the new authentication key:

```
security key-manager query
```

- Assign the new authentication key to your self-encrypting disks (SEDs):

```
storage encryption disk modify -disk disk_ID -data-key-id key_ID
```

Note: Make sure you are using the new authentication key from your query.

- If needed, assign a FIPS key to the SEDs:

```
storage encryption disk modify -disk disk_id -fips-key-id
fips_authentication_key_id
```

If your security setup requires you to use different keys for data authentication and FIPS 140-2 authentication, you should create a separate key for each. If that is not the case, you can use the same authentication key for FIPS compliance that you use for data access.

Related tasks

[Preparing to upgrade nodes using NetApp Storage Encryption with external key management servers](#) on page 31

Enabling and reverting LIFs to home ports (post-upgrade or downgrade tasks for the cluster)

During a reboot, some LIFs might have been migrated to their assigned failover ports. Before and after you upgrade, revert, or downgrade a cluster, you must enable and revert any LIFs that are not on their home ports.

About this task

The `network interface revert` command reverts a LIF that is not currently on its home port back to its home port, provided that the home port is operational. A LIF's home port is specified when the LIF is created; you can determine the home port for a LIF by using the `network interface show` command.

Steps

- Display the status of all LIFs:

```
network interface show
```

Example

This example displays the status of all LIFs for a storage virtual machine (SVM).

```
cluster1::> network interface show -vserver vs0
Logical      Status      Network      Current      Current      Is
Vserver      Interface   Admin/Oper   Address/Mask Node          Port         Home
-----
vs0
  data001     down/down   192.0.2.120/24  node0        e0e          true
  data002     down/down   192.0.2.121/24  node0        e0f          true
  data003     down/down   192.0.2.122/24  node0        e2a          true
  data004     down/down   192.0.2.123/24  node0        e2b          true
  data005     down/down   192.0.2.124/24  node0        e0e          false
  data006     down/down   192.0.2.125/24  node0        e0f          false
  data007     down/down   192.0.2.126/24  node0        e2a          false
  data008     down/down   192.0.2.127/24  node0        e2b          false
8 entries were displayed.
```

If any LIFs appear with a `Status Admin` status of `down` or with an `Is home` status of `false`, continue with the next step.

- Enable the data LIFs:

```
network interface modify {-role data} -status-admin up
```

Example

```
cluster1::> network interface modify {-role data} -status-admin up
8 entries were modified.
```

3. Revert LIFs to their home ports:

```
network interface revert *
```

Example

This command reverts all LIFs back to their home ports and changes all LIF home statuses to true.

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

4. Verify that all LIFs are in their home ports:

```
network interface show
```

Example

This example shows that all LIFs for SVM vs0 are on their home ports.

```
cluster1::> network interface show -vserver vs0
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
vs0	data001	up/up	192.0.2.120/24	node0	e0e	true
	data002	up/up	192.0.2.121/24	node0	e0f	true
	data003	up/up	192.0.2.122/24	node0	e2a	true
	data004	up/up	192.0.2.123/24	node0	e2b	true
	data005	up/up	192.0.2.124/24	node1	e0e	true
	data006	up/up	192.0.2.125/24	node1	e0f	true
	data007	up/up	192.0.2.126/24	node1	e2a	true
	data008	up/up	192.0.2.127/24	node1	e2b	true

8 entries were displayed.

Relocating moved load-sharing mirror source volumes

After successfully completing a nondisruptive upgrade, you can move load-sharing mirror source volumes back to the locations they were in originally before the upgrade.

Steps

1. Identify the location to which you are moving the load-sharing mirror source volume by using the record you created before moving the load-sharing mirror source volume.

Preparing all load-sharing mirrors for a major upgrade on page 27

2. Move the load-sharing mirror source volume back to its original location by using the `volume move start` command.

Resuming SnapMirror operations

After completing a nondisruptive upgrade or downgrade, you must resume any SnapMirror relationships that were suspended.

Before you begin

Existing SnapMirror relationships must have been suspended by using the `snapmirror quiesce` command, and the cluster must have been nondisruptively upgraded or downgraded.

Steps

1. Resume transfers for each SnapMirror relationship that was previously quiesced:

```
snapmirror resume *
```

This command resumes the transfers for all quiesced SnapMirror relationships.

2. Verify that the SnapMirror operations have resumed:

```
snapmirror show
```

Example

```
cluster1::> snapmirror show
```

Source Path	Type	Destination Path	Mirror State	Relationship Status	Total Progress	Healthy	Last Updated
cluster1-vs1:dp_src1	DP	cluster1-vs2:dp_dst1	Snapmirrored	Idle	-	true	-
cluster1-vs1:xdp_src1	XDP	cluster1-vs2:xdp_dst1	Snapmirrored	Idle	-	true	-
cluster1://cluster1-vs1/ls_src1	LS	cluster1://cluster1-vs1/ls_mr1	Snapmirrored	Idle	-	true	-
		cluster1://cluster1-vs1/ls_mr2	Snapmirrored	Idle	-	true	-

4 entries were displayed.

For each SnapMirror relationship, verify that the Relationship Status is "Idle". If the status is "Transferring", wait for the SnapMirror transfer to complete, and then reenter the command to verify that the status has changed to "Idle".

After you finish

For each SnapMirror relationship that is configured to run on a schedule, you should verify that the first scheduled SnapMirror transfer completes successfully.

Setting the desired NT ACL permissions display level for NFS clients

After upgrading from ONTAP 8.3.0, the default handling for displaying NT ACL permissions to NFS clients has changed. You should check the setting and change it to the desired setting for your environment if necessary. This task does not apply if you are upgrading from ONTAP 8.3.1 or later.

About this task

In multiprotocol environments, ONTAP displays to NFS clients the permissions of NTFS security-style files and directories based on the access granted by the NT ACL to any user. In ONTAP 8.3.0, ONTAP by default displayed to NFS clients the permission based on the maximum access granted by the NT ACL. After upgrading, the default setting changes to display permissions based on the minimum access granted by the NT ACL. This change applies to new and existing storage virtual machines (SVMs).

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Check the setting for displaying NT ACL permissions for NFS clients:

```
vserver nfs show -vserver vserver_name -fields ntacl-display-permissive-perms
```

After upgrading from 8.3.0, the value for this new parameter is disabled, meaning ONTAP displays the minimum permissions.

3. If you prefer to display the maximum permissions, change the setting individually for each SVM as desired:

```
vserver nfs modify -vserver vserver_name -ntacl-display-permissive-perms
enabled
```

4. Verify that the change took effect:

```
vserver nfs show -vserver vserver_name -fields ntacl-display-permissive-
perms
```

5. Return to the admin privilege level:

```
set -privilege admin
```

Enforcing SHA-2 on administrator account passwords

Administrator accounts created prior to ONTAP 9.0 continue to use MD5 passwords after the upgrade, until the passwords are manually changed. MD5 is less secure than SHA-2. Therefore, after upgrading, you should prompt users of MD5 accounts to change their passwords to use the default SHA-512 hash function.

About this task

The password hash functionality enables you to do the following:

- Display user accounts that match the specified hash function.
- Expire accounts that use a specified hash function (for example, MD5), forcing the users to change their passwords in their next login.
- Lock accounts whose passwords use the specified hash function.
- When reverting to a release earlier than ONTAP 9, reset the cluster administrator's own password for it to be compatible with the hash function (MD5) that is supported by the earlier release.

ONTAP accepts pre-hashed SHA-2 passwords only by using NetApp Manageability SDK (security-login-create and security-login-modify-password).

Manageability enhancements

Steps

1. Migrate the MD5 administrator accounts to the SHA-512 password hash function:

- a. Expire all MD5 administrator accounts:

```
security login expire-password -vserver * -username * -hash-function
md5
```

Doing so forces MD5 account users to change their passwords upon next login.

- b. Ask users of MD5 accounts to log in through a console or SSH session.

The system detects that the accounts are expired and prompts users to change their passwords. SHA-512 is used by default for the changed passwords.

2. Optional: For MD5 accounts whose users do not log in to change their passwords within a period of time, force the account migration:

- a. Lock accounts that still use the MD5 hash function (advanced privilege level):

```
security login expire-password -vserver * -username * -hash-function
md5 -lock-after integer
```

After the number of days specified by `-lock-after`, users cannot access their MD5 accounts.

- b. Unlock the accounts when the users are ready to change their passwords:

```
security login unlock -vserver vserver_name -username user_name
```

- c. Have users log in to their accounts through a console or SSH session and change their passwords when the system prompts them to do so.

When you need to update the Disk Qualification Package

The Disk Qualification Package (DQP) adds full support for newly qualified drives. Before you update drive firmware or add new drive types or sizes to a cluster, you must update the DQP. A best practice is to also update the DQP regularly; for example, every quarter or semi-annually.

You need to download and install the DQP in the following situations:

- Whenever you add a new drive type or size to the node
For example, if you already have 1-TB drives and add 2-TB drives, you need to check for the latest DQP update.
- Whenever you update the disk firmware
- Whenever newer disk firmware or DQP files are available
- Whenever you upgrade to a new version of ONTAP.
The DQP is not updated as part of an ONTAP upgrade.

Related information

[NetApp Downloads: Disk Qualification Package](#)

[NetApp Downloads: Disk Drive and Firmware](#)

Reverting clusters to an earlier ONTAP release

In some cases, to transition a cluster to an earlier ONTAP release, you must perform a *reversion*. Reverting is always disruptive, and it requires planning, preparation, the reversion itself, and several post-reversion procedures.

Attention: Do not attempt to revert ONTAP by simply downloading and booting (or netbooting) in an earlier release. If you do, you cannot boot the earlier target release. You must use the `clustershell system node revert-to` and `nodeshell revert_to` commands for the reversion process.

Related concepts

[Cluster update requirements](#) on page 10

When to revert and when to call technical support

You can downgrade or revert without assistance when downgrading or reverting new or test clusters, but you should call technical support if you encounter problems during or after upgrade, or if you want to downgrade or revert a production cluster.

You can revert or downgrade to an allowed ONTAP release without assistance from technical support only in the following scenarios:

- You upgraded to a new release on a test cluster and you want to return to the original release when testing is completed.
- You are configuring a new cluster—running a later release of ONTAP and not yet in production—in an environment in which you have standardized on an earlier ONTAP release.

If the upgrade fails, *do not* attempt to revert ONTAP in a production environment without assistance. If you encounter any of the following circumstances, contact technical support immediately:

- The upgrade process fails and cannot finish.
- The upgrade process finishes, but the cluster is unusable in a production environment.
- The upgrade process finishes and the cluster goes into production, but you are not satisfied with its behavior.
- The upgrade process finishes for some but not all of the nodes, and you decide that you want to revert.

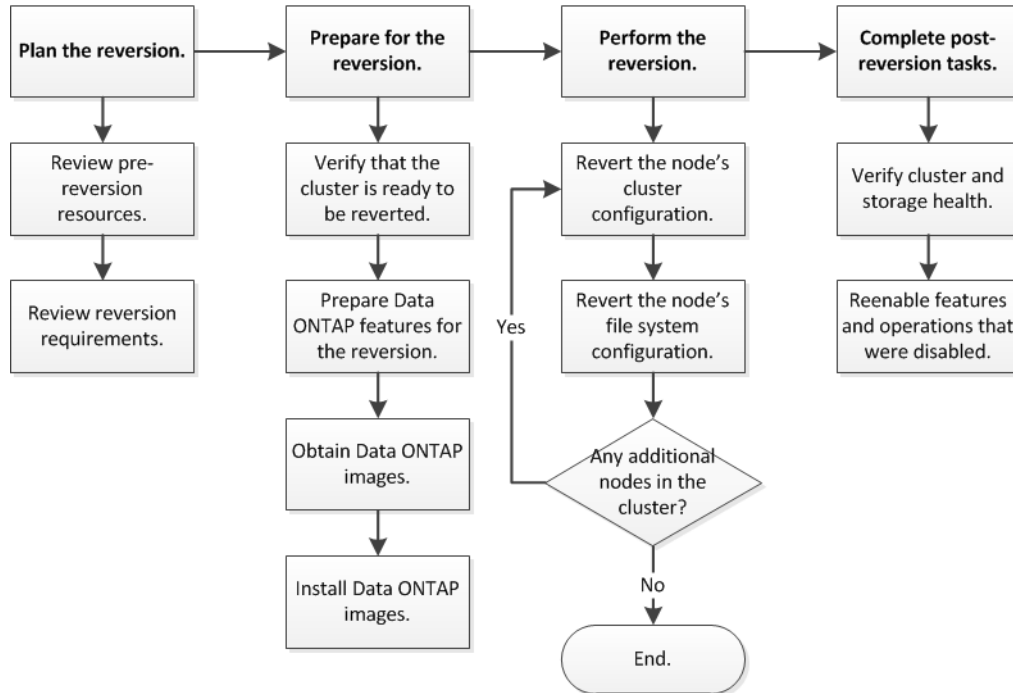
If you created volumes in ONTAP 9.5 or later and you need to revert to an earlier version, contact technical support to confirm if any of the volumes use adaptive compression. Volumes using adaptive compression must be uncompressed before reverting.

Related concepts

[Cluster update requirements](#) on page 10

Cluster revert workflow

You can use the cluster revert workflow to plan the reversion, prepare for the reversion, perform the reversion, and complete post-reversion tasks.



Planning your reversion

Because new features are introduced in each release of ONTAP, you must understand reversion requirements and evaluate how they might impact your current configuration.

Steps

1. [Reviewing pre-reversion resources](#) on page 85
Before reverting ONTAP, you should review resources to understand issues you must resolve before upgrading, understand new system behavior in the target release, and confirm hardware support.
2. [Reviewing cluster reversion requirements](#) on page 86
Before reverting ONTAP, you must verify that your cluster meets the general reversion requirements. Some configurations and features also have requirements that you should understand.

Reviewing pre-reversion resources

Before reverting ONTAP, you should review resources to understand issues you must resolve before upgrading, understand new system behavior in the target release, and confirm hardware support.

Steps

1. Review the *Release Notes* for the target release.

[ONTAP 9 Release Notes](#)

The “Important cautions” section describes potential issues that you should be aware of before upgrading to the new release. The “New and changed features” and “Known problems and limitations” sections describe new system behavior after upgrading to the new release.

2. Confirm that your hardware platform is supported in the target release.

NetApp Hardware Universe

3. Confirm that your cluster and management switches are supported in the target release.

You must verify that the NX-OS (cluster network switches), IOS (management network switches), and reference configuration file (RCF) software versions are compatible with the version of ONTAP to which you are reverting.

NetApp Downloads: Cisco Ethernet Switch

4. If your cluster is configured for SAN, confirm that the SAN configuration is fully supported.

All SAN components—including target ONTAP software version, host OS and patches, required Host Utilities software, and adapter drivers and firmware—should be supported.

NetApp Interoperability Matrix Tool

Reviewing cluster reversion requirements

Before reverting ONTAP, you must verify that your cluster meets the general reversion requirements. Some configurations and features also have requirements that you should understand.

Reversion process considerations

You need to consider the revert issues and limitations before beginning an ONTAP reversion.

- Reversion is disruptive.
No client access can occur during the reversion. If you are reverting a production cluster, be sure to include this disruption in your planning.
- Reversion affects all nodes in the cluster.
The reversion affects all nodes in the cluster; however, the reversion must be performed and completed on each HA pair before other HA pairs are reverted.
- The reversion is complete when all nodes are running the new target release.
When the cluster is in a mixed-version state, you should not enter any commands that alter the cluster operation or configuration except as necessary to satisfy reversion requirements; monitoring operations are permitted.
Attention: If you cannot complete the reversion for any reason, contact technical support immediately. If you have reverted some, but not all of the nodes, do not attempt to upgrade the cluster back to the source release.
- When you revert a node, it clears the cached data in a Flash Cache module.
Because there is no cached data in the Flash Cache module, the node serves initial read requests from disk, which results in decreased read performance during this period. The node repopulates the cache as it serves read requests.
- A LUN that is backed up to tape running on ONTAP 9.x can be restored only to 9.x and later releases and not to an earlier release.
- If your current version of ONTAP supports In-Band ACP (IBACP) functionality, and you revert to a version of ONTAP that does not support IBACP, the alternate path to your disk shelf is disabled.
- If LDAP is used by any of your storage virtual machines (SVMs), LDAP referral must be disabled before reversion.

Related concepts

Cluster update requirements on page 10

Reversion requirements for SnapMirror and SnapVault relationships

The `system node revert-to` command notifies you of any SnapMirror and SnapVault relationships that need to be deleted or reconfigured for the reversion process to be completed. However, you should be aware of these requirements before you begin the reversion.

- All SnapVault and data protection mirror relationships must be quiesced and then broken. After the reversion is completed, you can resynchronize and resume these relationships if a common Snapshot copy exists.
- SnapVault relationships must not contain the following SnapMirror policy types:
 - **async-mirror**
You must delete any relationship that uses this policy type.
 - **MirrorAndVault**
If any of these relationships exist, you should change the SnapMirror policy to **mirror-vault**.
- All load-sharing mirror relationships and destination volumes must be deleted.
- SnapMirror relationships with FlexClone destination volumes must be deleted.
- Network compression must be disabled for each SnapMirror policy.
- The `all_source_snapshot` rule must be removed from any **async-mirror** type SnapMirror policies.

Note: The Single File Snapshot Restore (SFSR) and Partial File Snapshot Restore (PFSR) operations are deprecated on the root volume.
- Any currently running single file and Snapshot restore operations must be completed before the reversion can proceed.
You can either wait for the restore operation to finish, or you can abort it.
- Any incomplete single file and Snapshot restore operations must be removed by using the `snapmirror restore` command.

Setting autocommit periods for SnapLock volumes before reverting

To revert from ONTAP 9, the value of the autocommit period for SnapLock volumes must be set in hours, not days. Before attempting to revert, you must check the autocommit value for your SnapLock volumes and modify it from days to hours, if necessary.

Steps

1. Verify that there are SnapLock volumes in the cluster that have unsupported autocommit periods:


```
volume snaplock show -autocommit-period *days
```
2. Modify the unsupported autocommit periods to hours:


```
volume snaplock modify -vserver vservers_name -volume volume_name -autocommit-period value hours
```

Preparing to revert ONTAP clusters

Before reverting to an earlier version of ONTAP, you must verify that the cluster is ready to be reverted and make any required configuration changes.

Steps

1. [Verifying that the cluster is ready to be reverted](#) on page 88
Before you perform the reversion, you should verify that your cluster configuration is healthy.
2. [Preparing to revert production clusters](#) on page 91
If you are reverting a cluster that you have configured to serve data to clients in your environment, you must ensure that certain configurations are prepared for the reversion.
3. [Obtaining ONTAP software images](#) on page 97
For ONTAP 9.4 and later, you can copy the ONTAP software image from the NetApp Support Site to a local folder. For upgrades from ONTAP 9.3 or earlier, you must copy the ONTAP software image to an HTTP server or FTP server on your network.

Verifying that the cluster is ready to be reverted

Before you perform the reversion, you should verify that your cluster configuration is healthy.

Verifying cluster health (verifying that the cluster is ready to be reverted)

Before and after you upgrade, revert, or downgrade a cluster, you should verify that the nodes are healthy and eligible to participate in the cluster, and that the cluster is in quorum.

Steps

1. Verify that the nodes in the cluster are online and are eligible to participate in the cluster:

```
cluster show
```

Example

```
cluster1::> cluster show
Node           Health  Eligibility
-----
node0          true   true
node1          true   true
```

If any node is unhealthy or ineligible, check EMS logs for errors and take corrective action.

2. Set the privilege level to advanced:
`set -privilege advanced`
3. Enter `y` to continue.
4. Verify the configuration details for each RDB process.
 - The relational database epoch and database epochs should match for each node.
 - The per-ring quorum master should be the same for all nodes.
Note that each ring might have a different quorum master.

To display this RDB process...	Enter this command...
Management application	<code>cluster ring show -unitname mgmt</code>
Volume location database	<code>cluster ring show -unitname vldb</code>
Virtual-Interface manager	<code>cluster ring show -unitname vifmgr</code>
SAN management daemon	<code>cluster ring show -unitname bcomd</code>

Example

This example shows the volume location database process:

```
cluster1::*> cluster ring show -unitname vldb
Node      UnitName Epoch    DB Epoch DB Trnxs Master  Online
-----
node0     vldb     154      154    14847 node0   master
node1     vldb     154      154    14847 node0   secondary
node2     vldb     154      154    14847 node0   secondary
node3     vldb     154      154    14847 node0   secondary
4 entries were displayed.
```

5. If you are operating in a SAN environment, verify that each node is in a SAN quorum:

```
event log show -messagename scsiblade.*
```

The most recent `scsiblade` event message for each node should indicate that the scsi-blade is in quorum.

Example

```
cluster1::*> event log show -messagename scsiblade.*
Time      Node      Severity  Event
-----
MM/DD/YYYY TIME node0     INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
MM/DD/YYYY TIME node1     INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
```

6. Return to the admin privilege level:

```
set -privilege admin
```

Related information

[System administration](#)

Verifying storage health (verifying that the cluster is ready to be reverted)

Before and after you upgrade, revert, or downgrade a cluster, you should verify the status of your disks, aggregates, and volumes.

Steps

1. If you are preparing to upgrade, revert, or downgrade, verify disk status:

To check for...	Do this...
Broken disks	<ol style="list-style-type: none"> Display any broken disks: <code>storage disk show -state broken</code> Remove or replace any broken disks.
Disks undergoing maintenance or reconstruction	<ol style="list-style-type: none"> Display any disks in maintenance, pending, or reconstructing states: <code>storage disk show -state maintenance pending reconstructing</code> Wait for the maintenance or reconstruction operation to finish before proceeding.

2. Verify that all aggregates are online by displaying the state of physical and logical storage, including storage aggregates:

```
storage aggregate show -state !online
```

This command displays the aggregates that are *not* online. All aggregates must be online before and after performing a major upgrade or reversion.

Example

```
cluster1::> storage aggregate show -state !online
There are no entries matching your query.
```

3. Verify that all volumes are online by displaying any volumes that are *not* online:

```
volume show -state !online
```

All volumes must be online before and after performing a major upgrade or reversion.

Example

```
cluster1::> volume show -state !online
There are no entries matching your query.
```

4. Verify that there are no inconsistent volumes:

```
volume show -is-inconsistent true
```

If any inconsistent volumes are returned, you must contact NetApp Support before you proceed with the upgrade.

Related information

[Disk and aggregate management](#)

Verifying the system time

You should verify that NTP is configured, and that the time is synchronized across the cluster.

Steps

1. Verify that the cluster is associated with an NTP server:

```
cluster time-service ntp server show
```

2. Verify that each node has the same date and time:

```
cluster date show
```

Example

```
cluster1::> cluster date show
Node      Date              Timezone
-----
node0     4/6/2013 20:54:38 GMT
node1     4/6/2013 20:54:38 GMT
node2     4/6/2013 20:54:38 GMT
node3     4/6/2013 20:54:38 GMT
4 entries were displayed.
```

Preparing to revert production clusters

If you are reverting a cluster that you have configured to serve data to clients in your environment, you must ensure that certain configurations are prepared for the reversion.

Considerations for reverting systems with SnapMirror Synchronous relationships

You must be aware of the considerations for SnapMirror Synchronous relationships before reverting from ONTAP 9.6 to ONTAP 9.5.

Before reverting, you must take the following steps if you have SnapMirror Synchronous relationships:

- You must delete any SnapMirror Synchronous relationship in which the source volume is serving data using NFSv4 or SMB/CIFS.
ONTAP 9.5 does not support NFSv4 and SMB/CIFS.
- You must delete any SnapMirror Synchronous relationships in a mirror-mirror cascade deployment.
A mirror-mirror cascade deployment is not supported for SnapMirror Synchronous relationships in ONTAP 9.5.
- If the common Snapshot copies in ONTAP 9.5 are not available during revert, you must initialize the SnapMirror Synchronous relationship after reverting.
After two hours of upgrade to ONTAP 9.6, the common Snapshot copies from ONTAP 9.5 are automatically replaced by the common Snapshot copies in ONTAP 9.6. Therefore, you cannot resynchronize the SnapMirror Synchronous relationship after reverting if the common Snapshot copies from ONTAP 9.5 are not available.

Reversing physical block sharing in split FlexClone volumes

If you have split a FlexClone volume from its parent volume, you must undo the sharing of any physical block between the clone and its parent volume before reverting from ONTAP 9.4 or later to an earlier version of ONTAP.

About this task

This task is applicable only for AFF systems when split has been run on any of the FlexClone volumes.

Steps

1. Log in to the advanced privilege level:
`set -privilege advanced`
2. Identify the split FlexClone volumes with shared physical blocks:
`volume clone sharing-by-split show`

Example

```
cluster1::> volume clone sharing-by-split show
Node          Vserver  Volume      Aggregate
-----
node1         vs1      vol_clone1  aggr1
node2         vs2      vol_clone2  aggr2
2 entries were displayed.
```

3. Undo the physical block sharing in all of the split FlexClone volumes across the cluster:
`volume clone sharing-by-split undo start-all`

- Verify that there are no split FlexClone volumes with shared physical blocks:

```
volume clone sharing-by-split show
```

Example

```
cluster1::> volume clone sharing-by-split show
This table is currently empty.
```

Disabling qtree functionality in FlexGroup volumes before reverting to an earlier version of ONTAP

Qtrees for FlexGroup volumes are not supported prior to ONTAP 9.3. You must disable the qtree functionality on FlexGroup volumes before reverting from ONTAP 9.3 to an earlier version of ONTAP.

About this task

The qtree functionality is enabled either when you create a qtree or if you modify the `security-style` and `oplock-mode` attributes of the default qtree.

Steps

- Identify and delete all of the non-default qtrees in each FlexGroup volume that are enabled with the qtree functionality:
 - Log in to the advanced privilege level:


```
set -privilege advanced
```
 - Verify if any FlexGroup volume is enabled with the qtree functionality.

Example

For ONTAP 9.6 or later, use:

```
volume show is-qtrees-enabled true
```

For ONTAP 9.5 or earlier, use:

```
volume show -is-flexgroup-qtrees-enabled true
```

```
cluster1::*> volume show -is-flexgroup-qtrees-enabled true
Vserver  Volume      Aggregate  State  Type  Size  Available  Used%
-----  -
vs0      fg           -          online RW    320MB  220.4MB  31%
```

- Delete all of the non-default qtrees in each FlexGroup volume that are enabled with the qtree functionality:

```
volume qtree delete -vserver svm_name -volume volume_name -qtree
qtree_name
```

If the qtree functionality is enabled because you modified the attributes of the default qtree and if you do not have any qtrees, you can skip this step.

Example

```
cluster1::*> volume qtree delete -vserver vs0 -volume fg -qtree qtree4
WARNING: Are you sure you want to delete qtree qtree4 in volume fg vserver
vs0? {y|n}: y
[Job 38] Job is queued: Delete qtree qtree4 in volume fg vserver vs0.
```

- Disable the qtree functionality on each FlexGroup volume:

```
volume flexgroup qtree-disable -vserver svm_name -volume volume_name
```

Example

```
cluster1::*> volume flexgroup qtree-disable -vserver vs0 -volume fg
```

3. Identify and delete any Snapshot copies that are enabled with the qtree functionality.
 - a. Verify if any Snapshot copies are enabled with the qtree functionality:

```
volume snapshot show -vserver vserver_name -volume volume_name -fields is-flexgroup-qtree-enabled
```

Example

```
cluster1::*> volume snapshot show -vserver vs0 -volume fg -fields is-
flexgroup-qtree-enabled
vserver volume snapshot is-flexgroup-qtree-enabled
-----
vs0      fg      fg_snap1 true
vs0      fg      daily.2017-09-27_0010 true
vs0      fg      daily.2017-09-28_0010 true
vs0      fg      snapmirror.0241f354-a865-11e7-
alc0-00a098a71764_2147867740.2017-10-04_124524 true
```

- b. Delete all of the Snapshot copies that are enabled with the qtree functionality:

```
volume snapshot delete -vserver svm_name -volume volume_name -snapshot
snapshot_name -force true -ignore-owners true
```

The Snapshot copies that must be deleted include regular Snapshot copies and the Snapshot copies taken for SnapMirror relationships. If you have created any SnapMirror relationship for the FlexGroup volumes with a destination cluster that is running ONTAP 9.2 or earlier, you must delete all of the Snapshot copies that were taken when the source FlexGroup volume was enabled for the qtree functionality.

Example

```
cluster1:::> volume snapshot delete -vserver vs0 -volume fg -snapshot daily.
2017-09-27_0010 -force true -ignore-owners true
```

Related information

[FlexGroup volumes management](#)

Identifying and moving CIFS servers in workgroup mode

Before performing a revert, you must delete any CIFS servers in workgroup mode or move them in to a domain. Workgroup mode is not supported on ONTAP versions prior to ONTAP 9.

Steps

1. Identify any CIFS servers with a Authentication Style of workgroup:

```
vserver cifs show
```

2. Move or delete the servers you identified:

If you are going to...	Then use this command...
Move the CIFS server from the workgroup to an Active Directory domain:	<pre>vserver cifs modify -vserver vserver_name -domain domain_name</pre>

If you are going to...	Then use this command...
Delete the CIFS server	<code>vserver cifs delete -vserver vserver_name</code>

3. If you deleted the CIFS server, enter the username of the domain, then enter the user password.

Related concepts

[Cluster update requirements](#) on page 10

Related information

[SMB/CIFS management](#)

Reverting systems with deduplicated volumes

Before reverting from any version of ONTAP 9, you must ensure that the volumes contain sufficient free space for the revert operation.

Before you begin

The volume must have enough space to accommodate the savings that were achieved through the inline detection of blocks of zeros. For information about the space required, contact technical support.

About this task

Reverting from ONTAP 9 on a system that has deduplication enabled includes running advanced mode commands. You must contact technical support for assistance.

If you have enabled both deduplication and data compression on a volume that you want to revert, then you must revert data compression before reverting deduplication.

Steps

1. Use the `volume efficiency show` command with the `-fields` option to view the progress of the efficiency operations that are running on the volumes.

Example

The following command displays the progress of efficiency operations:

```
volume efficiency show -fields vserver,volume,progress
```

2. Use the `volume efficiency stop` command with the `-all` option to stop all active and queued deduplication operations.

Example

The following command stops all active and queued deduplication operations on volume VolA:

```
volume efficiency stop -vserver vs1 -volume VolA -all
```

3. Use the `set -privilege advanced` command to log in at the advanced privilege level.
4. Use the `volume efficiency revert-to` command with the `-version` option to downgrade the efficiency metadata of a volume to a specific version of ONTAP.

Example

The following command reverts the efficiency metadata on volume VolA to ONTAP 9.x:

```
volume efficiency revert-to -vserver vs1 -volume VolA -version 9.x
```

Note: The `volume efficiency revert-to` command reverts volumes that are present on the node on which this command is executed. This command does not revert volumes across nodes.

5. Use the `volume efficiency show` command with the `-op-status` option to monitor the progress of the downgrade.

Example

The following command monitors and displays the status of the downgrade:

```
volume efficiency show -vserver vs1 -op-status Downgrading
```

6. If the revert does not succeed, use the `volume efficiency show` command with the `-instance` option to see why the revert failed.

Example

The following command displays detailed information about all fields:

```
volume efficiency show -vserver vs1 -volume vol1 - instance
```

7. After the revert operation is complete, return to the admin privilege level:

```
set -privilege admin
```

Logical storage management

Reverting two-node and four-node MetroCluster configurations

Before reverting a two-node or four-node MetroCluster configuration, you must disable automatic unplanned switchover (AUSO).

Step

1. On both the clusters in MetroCluster, disable automatic unplanned switchover:


```
metrocluster modify -auto-switchover-failure-domain auso-disabled
```

Related information

MetroCluster management and disaster recovery

Preparing Snapshot copies before reverting

Before reverting to an earlier ONTAP release, you must disable all Snapshot copy policies and delete any Snapshot copies that were created after upgrading to the current release.

Before you begin

If you are reverting in a SnapMirror environment, you must first have deleted the following mirror relationships:

- All load-sharing mirror relationships
- Any data protection mirror relationships that were created in ONTAP 8.3.x
- All data protection mirror relationships if the cluster was re-created in ONTAP 8.3.x

Steps

1. Disable Snapshot copy policies for all data SVMs:


```
volume snapshot policy modify -vserver * -enabled false
```

2. Disable Snapshot copy policies for each node's aggregates:
 - a. Identify the node's aggregates by using the run `-node nodename aggr status` command.
 - b. Disable the Snapshot copy policy for each aggregate:


```
run -node nodename aggr options aggr_name nosnap on
```
 - c. Repeat this step for each remaining node.
3. Disable Snapshot copy policies for each node's root volume:
 - a. Identify the node's root volume by using the run `-node nodename vol status` command.
You identify the root volume by the word `root` in the `Options` column of the `vol status` command output.

Example

```
vsl::> run -node node1 vol status
```

Volume	State	Status	Options
vol0	online	raid_dp, flex 64-bit	root, nvfail=on

- b. Disable the Snapshot copy policy on the root volume:


```
run -node nodename vol options root_volume_name nosnap on
```
 - c. Repeat this step for each remaining node.
4. Delete all Snapshot copies that were created after upgrading to the current release:
 - a. Set the privilege level to advanced:


```
set -privilege advanced
```
 - b. Disable the snapshots:


```
snapshot policy modify -vserver * -enabled false
```
 - c. Delete the node's newer-version Snapshot copies:


```
volume snapshot prepare-for-revert -node nodename
```

This command deletes the newer-version Snapshot copies on each data volume, root aggregate, and root volume.

If any Snapshot copies cannot be deleted, the command fails and notifies you of any required actions you must take before the Snapshot copies can be deleted. You must complete the required actions and then rerun the `volume snapshot prepare-for-revert` command before proceeding to the next step.

Example

```
cluster1::*> volume snapshot prepare-for-revert -node node1
```

Warning: This command will delete all Snapshot copies that have the format used by the current version of ONTAP. It will fail if any Snapshot copy policies are enabled, or if any Snapshot copies have an owner. Continue? {y|n}: y

- d. Verify that the Snapshot copies have been deleted:

```
volume snapshot show -node nodename
```

If any newer-version Snapshot copies remain, force them to be deleted:


```
volume snapshot delete {-fs-version 9.0 -node nodename -is-constituent true} -ignore-owners -force
```

e. Repeat this step c for each remaining node.

f. Return to the admin privilege level:

```
set -privilege admin
```

Note: You must perform these steps on both the clusters in MetroCluster configuration.

Setting autocommit periods for SnapLock volumes before reverting

To revert from ONTAP 9, the value of the autocommit period for SnapLock volumes must be set in hours, not days. Before attempting to revert, you must check the autocommit value for your SnapLock volumes and modify it from days to hours, if necessary.

Steps

1. Verify that there are SnapLock volumes in the cluster that have unsupported autocommit periods:

```
volume snaplock show -autocommit-period *days
```

2. Modify the unsupported autocommit periods to hours:

```
volume snaplock modify -vserver vservers_name -volume volume_name -autocommit-period value hours
```

Obtaining ONTAP software images

For ONTAP 9.4 and later, you can copy the ONTAP software image from the NetApp Support Site to a local folder. For upgrades from ONTAP 9.3 or earlier, you must copy the ONTAP software image to an HTTP server or FTP server on your network.

About this task

To upgrade, revert, or downgrade the cluster to the target release of ONTAP, you require access to software images. Software images, firmware version information, and the latest firmware for your platform model are available on the NetApp Support Site. You should note the following important information:

- Software images are specific to platform models. You must obtain the correct image for your cluster.
- Software images include the latest version of system firmware that was available when a given version of ONTAP was released.
- If you are upgrading a system with NetApp Volume Encryption to ONTAP 9.5 or later, you must download the ONTAP software image for non-restricted countries, which includes NetApp Volume Encryption.

If you use the ONTAP software image for restricted countries to upgrade a system with NetApp Volume Encryption, the system panics and you lose access to your volumes.

Steps

1. Locate the target ONTAP software in the **Software Downloads** area of the NetApp Support Site.
2. Copy the software image.
 - For ONTAP 9.3 or earlier, copy the software image (for example, `93_q_image.tgz`) from the NetApp Support Site to the directory on the HTTP server or FTP server from which the image will be served

- For ONTAP 9.4 or later, copy the software image (for example, `95_q_image.tgz`) from the NetApp Support Site to the directory on the HTTP server or FTP server from which the image will be served or to a local folder.

Related information

[NetApp Downloads: Software](#)

Installing ONTAP software images for a reversion

Before performing a reversion, you must install the target ONTAP software image on each node in the cluster.

Before you begin

You must have obtained the ONTAP software images.

Steps

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

```
set -privilege advanced
```

The advanced prompt (`*>`) appears.

- Choose one of the following options based on your requirements:

If you want to...	Run this command...
Download, but not install, the software image	<pre>system node image get -node * -package location -replace-package true -background true</pre> <p>This command downloads the software image to all of the nodes simultaneously. To download the image to each node one at a time, do not specify the <code>-background</code> parameter.</p>
Install a previously downloaded software image	<pre>system node image update -node * -package image_name -background true</pre> <p>Note the following considerations for this command:</p> <ul style="list-style-type: none"> You need to set the privilege level to advanced (<code>set -privilege advanced</code>), entering <code>y</code> when prompted to continue. If you are unsure of the image name to install, then you can view a list of previously downloaded software images by using the <code>system node image package show</code> command. 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.
Download and install the software image in the same operation	<pre>system node image update -node * -package location -replace-package true -background true</pre> <p>Note the following considerations for this command:</p> <ul style="list-style-type: none"> You need to set the privilege level to advanced (<code>set -privilege advanced</code>), entering <code>y</code> when prompted to continue. This command downloads and installs the software image on all of the nodes simultaneously. To download and install the image on each node one at a time, do not specify the <code>-background</code> parameter.

3. Verify that the software image is downloaded and installed on each node:

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

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

Example

The following example shows a 2-node cluster in which the software image has been downloaded and 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:  Installation complete. image2 updated on node node0.
There is no update/install in progress
Status of most recent operation:
  Run Status:    Exited
  Exit Status:   Success
  Phase:         Run Script
  Exit Message:  Installation complete. image2 updated on node node1.
2 entries were acted on.
```

Reverting an ONTAP cluster

To take the cluster offline to revert to an earlier ONTAP release, you must disable storage failover and the data LIFs, address reversion preconditions, revert the cluster and file system configurations on a node, and then repeat the process for each additional node in the cluster.

Before you begin

You must have satisfied reversion preparation requirements.

About this task

Reverting a cluster requires you to take the cluster offline for the duration of the reversion.

Steps

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

```
set -privilege advanced
```

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

2. Verify that the target ONTAP software is installed:

```
system image show
```

Example

The following example shows that version 9.1 is installed as the alternate image on both nodes:

```
cluster1:*> system image show
Node      Image  Is      Is      Install
          Image Default Current Version Date
-----
node0
  image1  true   true   9.2    MM/DD/YYYY TIME
  image2  false  false  9.1    MM/DD/YYYY TIME
```

```
node1
  image1 true true 9.2 MM/DD/YYYY TIME
  image2 false false 9.1 MM/DD/YYYY TIME
4 entries were displayed.
```

3. Disable all of the data LIFs in the cluster:

```
network interface modify {-role data} -status-admin down
```

4. If the cluster consists of only two nodes, disable cluster HA:

```
cluster ha modify -configured false
```

5. Disable storage failover for the nodes in the HA pair from either node:

```
storage failover modify -node nodename -enabled false
```

You only need to disable storage failover once for the HA pair. When you disable storage failover for a node, storage failover is also disabled on the node's partner.

6. Log in to the node that you want to revert.

To revert a node, you must be logged in to the cluster through the node's node management LIF.

7. Set the node's target ONTAP software image to be the default image:

```
system image modify -node nodename -image target_image -isdefault true
```

8. Verify that the target ONTAP software image is set as the default image for the node that you are reverting:

```
system image show
```

Example

The following example shows that version 9.1 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 9.2 MM/DD/YYYY TIME
  image2 true false 9.1 MM/DD/YYYY TIME
node1
  image1 true true 9.2 MM/DD/YYYY TIME
  image2 false false 9.1 MM/DD/YYYY TIME
4 entries were displayed.
```

9. If the cluster consists of only two nodes, verify that the node does not hold epsilon:

- a. Check whether the node currently holds epsilon:

```
cluster show -node nodename
```

Example

The following example shows that the node holds epsilon:

```
cluster1:*> cluster show -node node1
Node: node1
UUID: 026efc12-ac1a-11e0-80ed-0f7eba8fc313
Epsilon: true
Eligibility: true
Health: true
```

- b. If the node holds epsilon, mark epsilon as **false** on the node so that epsilon can be transferred to the node's partner:

```
cluster modify -node nodenameA -epsilon false
```

- c. Transfer epsilon to the node's partner by marking epsilon `true` on the partner node:

```
cluster modify -node nodenameB -epsilon true
```

10. Verify that the node is ready for reversion:

```
system node revert-to -node nodename -check-only true -version 9.x
```

The `check-only` parameter identifies any preconditions that must be addressed before reverting, such as the following examples:

- Disabling storage failover
- Disabling the Snapshot policy
- Deleting Snapshot copies that were created after upgrading to the later version of ONTAP

11. Verify that all of the preconditions have been addressed:

```
system node revert-to -node nodename -check-only true -version 9.x
```

12. Revert the cluster configuration of the node:

```
system node revert-to -node nodename -version 9.x
```

The `-version` option refers to the target release. For example, if the software you installed and verified is ONTAP 9.1, the correct value of the `-version` option is `9.1`.

The cluster configuration is reverted, and then you are logged out of the clustershell.

13. Log back in to the clustershell, and then switch to the nodeshell:

```
run -node nodename
```

After logging on the clustershell again, it might take a few minutes before it is ready to accept the nodeshell command. So, if the command fails, wait a few minutes and try it again.

14. Revert the file system configuration of the node:

```
revert_to 9.x
```

This command verifies that the node's file system configuration is ready to be reverted, and then reverts it. If any preconditions are identified, you must address them and then rerun the `revert_to` command.

Note: Using a system console to monitor the revert process displays greater details than seen in nodeshell.

When the command finishes, the `LOADER` prompt is displayed.

15. Enter `yes` at prompt to revert.

If `AUTOBOOT` is true, the node will reboot to ONTAP. If `AUTOBOOT` is false, the node will halt.

16. Repeat Steps 5 through 15 on the other node in the HA pair.

17. If the cluster consists of only two nodes, reenabling cluster HA:

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

18. Reenable storage failover on both nodes if it was previously disabled:

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

19. Repeat Steps 4 through 18 for each additional HA pair and both the clusters in MetroCluster Configuration.

Completing post-reversion tasks

After reverting to an earlier version of ONTAP, you might need to perform additional tasks to provide cluster health and storage availability.

Steps

1. [Enabling automatic switchover for MetroCluster configurations](#) on page 102
This topic provides information regarding the additional tasks that you must perform after the reversion of MetroCluster configurations.
2. [Verifying cluster health \(completing post-reversion tasks\)](#) on page 103
Before and after you upgrade, revert, or downgrade a cluster, you should verify that the nodes are healthy and eligible to participate in the cluster, and that the cluster is in quorum.
3. [Verifying storage health \(completing post-reversion tasks\)](#) on page 104
Before and after you upgrade, revert, or downgrade a cluster, you should verify the status of your disks, aggregates, and volumes.
4. [Enabling and reverting LIFs to home ports \(completing post-reversion tasks\)](#) on page 105
During a reboot, some LIFs might have been migrated to their assigned failover ports. Before and after you upgrade, revert, or downgrade a cluster, you must enable and revert any LIFs that are not on their home ports.
5. [Preparing Snapshot copies after reverting](#) on page 106
After reverting to an earlier version of ONTAP, you must enable Snapshot copy policies to start creating Snapshot copies again.
6. [Verifying client access \(CIFS and NFS\)](#) on page 106
For the configured protocols, test access from CIFS and NFS clients to verify that the cluster is accessible.
7. [Verifying IPv6 firewall entries](#) on page 107
A reversion from any version of ONTAP 9 might result in missing default IPv6 firewall entries for some services in firewall policies. You need to verify that the required firewall entries have been restored to your system.
8. [Reverting password hash function to the supported encryption type](#) on page 108
If you revert to a release prior from any version of ONTAP 9, SHA-2 account users can no longer be authenticated with their passwords. Therefore, you must have them reset their passwords to using the encryption type (MD5) that is supported by the release you revert to.
9. [Considerations for whether to manually update the SP firmware](#) on page 108
If the SP automatic update functionality is enabled (the default), downgrading or reverting to ONTAP 8.3.x does not require a manual SP firmware update. The SP firmware is automatically updated to the newest compatible version that is supported by the ONTAP version you reverted or downgraded to.

Enabling automatic switchover for MetroCluster configurations

This topic provides information regarding the additional tasks that you must perform after the reversion of MetroCluster configurations.

Steps

1. Enable automatic unplanned switchover:

```
metrocluster modify -auto-switchover-failure-domain auto-on-cluster-disaster
```
2. Validate the MetroCluster configuration:

```
metrocluster check run
```

Verifying cluster health (completing post-reversion tasks)

Before and after you upgrade, revert, or downgrade a cluster, you should verify that the nodes are healthy and eligible to participate in the cluster, and that the cluster is in quorum.

Steps

1. Verify that the nodes in the cluster are online and are eligible to participate in the cluster:

```
cluster show
```

Example

```
cluster1::> cluster show
Node           Health  Eligibility
-----
node0          true   true
node1          true   true
```

If any node is unhealthy or ineligible, check EMS logs for errors and take corrective action.

2. Set the privilege level to advanced:

```
set -privilege advanced
```

3. Enter **y** to continue.

4. Verify the configuration details for each RDB process.

- The relational database epoch and database epochs should match for each node.
- The per-ring quorum master should be the same for all nodes.
Note that each ring might have a different quorum master.

To display this RDB process...	Enter this command...
Management application	<code>cluster ring show -unitname mgmt</code>
Volume location database	<code>cluster ring show -unitname vldb</code>
Virtual-Interface manager	<code>cluster ring show -unitname vifmgr</code>
SAN management daemon	<code>cluster ring show -unitname bcomd</code>

Example

This example shows the volume location database process:

```
cluster1::*> cluster ring show -unitname vldb
Node      UnitName Epoch  DB Epoch DB Trnxs Master  Online
-----
node0     vldb     154    154     14847  node0  master
node1     vldb     154    154     14847  node0  secondary
node2     vldb     154    154     14847  node0  secondary
node3     vldb     154    154     14847  node0  secondary
4 entries were displayed.
```

5. If you are operating in a SAN environment, verify that each node is in a SAN quorum:

```
event log show -messagename scsiblade.*
```

The most recent `scsiblade` event message for each node should indicate that the scsi-blade is in quorum.

Example

```
cluster1::*> event log show -messagename scsiblade.*
Time           Node           Severity      Event
-----
MM/DD/YYYY TIME node0           INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
MM/DD/YYYY TIME node1           INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
```

- Return to the admin privilege level:

```
set -privilege admin
```

Related information

[System administration](#)

Verifying storage health (completing post-reversion tasks)

Before and after you upgrade, revert, or downgrade a cluster, you should verify the status of your disks, aggregates, and volumes.

Steps

- If you are preparing to upgrade, revert, or downgrade, verify disk status:

To check for..	Do this..
Broken disks	<ol style="list-style-type: none"> Display any broken disks: disk show -container-type broken Remove or replace any broken disks.
Disks undergoing maintenance or reconstruction	<ol style="list-style-type: none"> Display any disks in maintenance, pending, or reconstructing states: storage disk show -state maintenance pending reconstructing Wait for the maintenance or reconstruction operation to finish before proceeding.

- Verify that all aggregates are online by displaying the state of physical and logical storage, including storage aggregates:

```
storage aggregate show -state !online
```

This command displays the aggregates that are *not* online. All aggregates must be online before and after performing a major upgrade or reversion.

Example

```
cluster1::> storage aggregate show -state !online
There are no entries matching your query.
```

- Verify that all volumes are online by displaying any volumes that are *not* online:

```
volume show -state !online
```

All volumes must be online before and after performing a major upgrade or reversion.

Example

```
cluster1::> volume show -state !online
There are no entries matching your query.
```

4. Verify that there are no inconsistent volumes:

```
volume show -is-inconsistent true
```

If any inconsistent volumes are returned, you must contact NetApp Support before you proceed with the upgrade.

Related information

[Disk and aggregate management](#)

Enabling and reverting LIFs to home ports (completing post-reversion tasks)

During a reboot, some LIFs might have been migrated to their assigned failover ports. Before and after you upgrade, revert, or downgrade a cluster, you must enable and revert any LIFs that are not on their home ports.

About this task

The `network interface revert` command reverts a LIF that is not currently on its home port back to its home port, provided that the home port is operational. A LIF's home port is specified when the LIF is created; you can determine the home port for a LIF by using the `network interface show` command.

Steps

1. Display the status of all LIFs:

```
network interface show
```

Example

This example displays the status of all LIFs for a storage virtual machine (SVM).

```
cluster1::> network interface show -vserver vs0
-----
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
vs0	data001	down/down	192.0.2.120/24	node0	e0e	true
	data002	down/down	192.0.2.121/24	node0	e0f	true
	data003	down/down	192.0.2.122/24	node0	e2a	true
	data004	down/down	192.0.2.123/24	node0	e2b	true
	data005	down/down	192.0.2.124/24	node0	e0e	false
	data006	down/down	192.0.2.125/24	node0	e0f	false
	data007	down/down	192.0.2.126/24	node0	e2a	false
	data008	down/down	192.0.2.127/24	node0	e2b	false

```
-----
8 entries were displayed.
```

If any LIFs appear with a `Status Admin` status of down or with an `Is home` status of false, continue with the next step.

2. Enable the data LIFs:

```
network interface modify {-role data} -status-admin up
```

Example

```
cluster1::> network interface modify {-role data} -status-admin up
8 entries were modified.
```

3. Revert LIFs to their home ports:

```
network interface revert *
```

Example

This command reverts all LIFs back to their home ports and changes all LIF home statuses to true.

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

4. Verify that all LIFs are in their home ports:

```
network interface show
```

Example

This example shows that all LIFs for SVM vs0 are on their home ports.

```
cluster1::> network interface show -vserver vs0
-----
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
vs0	data001	up/up	192.0.2.120/24	node0	e0e	true
	data002	up/up	192.0.2.121/24	node0	e0f	true
	data003	up/up	192.0.2.122/24	node0	e2a	true
	data004	up/up	192.0.2.123/24	node0	e2b	true
	data005	up/up	192.0.2.124/24	node1	e0e	true
	data006	up/up	192.0.2.125/24	node1	e0f	true
	data007	up/up	192.0.2.126/24	node1	e2a	true
	data008	up/up	192.0.2.127/24	node1	e2b	true

```
-----
8 entries were displayed.
```

Preparing Snapshot copies after reverting

After reverting to an earlier version of ONTAP, you must enable Snapshot copy policies to start creating Snapshot copies again.

About this task

You are reenabling the Snapshot schedules that you disabled before you reverted to an earlier version of ONTAP.

Steps

1. Enable Snapshot copy policies for all data SVMs:

```
volume snapshot policy modify -vserver * -enabled true
snapshot policy modify pg-rop-hourly -enable true
```

2. For each node, enable the Snapshot copy policy of the root volume by using the run -node *nodename* vol options *root_vol_name* nosnap off command.

Example

```
cluster1::> run -node node1 vol options vol0 nosnap off
```

Verifying client access (CIFS and NFS)

For the configured protocols, test access from CIFS and NFS clients to verify that the cluster is accessible.

Verifying IPv6 firewall entries

A reversion from any version of ONTAP 9 might result in missing default IPv6 firewall entries for some services in firewall policies. You need to verify that the required firewall entries have been restored to your system.

Steps

1. Verify that all firewall policies are correct by comparing them to the default policies:

```
system services firewall policy show
```

Example

The following example shows the default policies:

```
cluster1::*> system services firewall policy show
Policy          Service      Action IP-List
-----
cluster
                dns         allow  0.0.0.0/0
                http        allow  0.0.0.0/0
                https       allow  0.0.0.0/0
                ndmp        allow  0.0.0.0/0
                ntp         allow  0.0.0.0/0
                rsh         allow  0.0.0.0/0
                snmp        allow  0.0.0.0/0
                ssh         allow  0.0.0.0/0
                telnet      allow  0.0.0.0/0
data
                dns         allow  0.0.0.0/0, ::/0
                http        deny   0.0.0.0/0, ::/0
                https       deny   0.0.0.0/0, ::/0
                ndmp        allow  0.0.0.0/0, ::/0
                ntp         deny   0.0.0.0/0, ::/0
                rsh         deny   0.0.0.0/0, ::/0
.
.
.
```

2. Manually add any missing default IPv6 firewall entries by creating a new firewall policy:

```
system services firewall policy create
```

Example

```
cluster1::*> system services firewall policy create -policy newIPv6
               -service ssh -action allow -ip-list ::/0
```

3. Apply the new policy to the LIF to allow access to a network service:

```
network interface modify
```

Example

```
cluster1::*> network interface modify -vserver VS1 -lif LIF1
               -firewall-policy newIPv6
```

Reverting password hash function to the supported encryption type

If you revert to a release prior from any version of ONTAP 9, SHA-2 account users can no longer be authenticated with their passwords. Therefore, you must have them reset their passwords to using the encryption type (MD5) that is supported by the release you revert to.

Steps

1. Prior to the revert, identify the user accounts that use the SHA-2 hash function (advanced privilege level):

```
security login show -vserver * -username * -application * -
authentication-method password -hash-function !md5
```

You should retain the command output. You need the account information after the revert.

2. During the revert, run the advanced command `security Login password-prepare-to-downgrade` as prompted to reset your own password to using the MD5 hash function.

If your password is not encrypted with MD5, the command prompts you for a new password and encrypts it with MD5, enabling your credential to be authenticated after the revert.

3. After the revert, reset SHA-2 accounts to MD5:

- a. For each SHA-2 account you identified, change the password to a temporary one:

```
security login password -username user_name -vserver vserver_name
```

The changed password uses the MD5 hash function.

- b. Communicate the temporary password to the affected users and have them log in through a console or SSH session to change their passwords as prompted by the system.

Considerations for whether to manually update the SP firmware

If the SP automatic update functionality is enabled (the default), downgrading or reverting to ONTAP 8.3.x does not require a manual SP firmware update. The SP firmware is automatically updated to the newest compatible version that is supported by the ONTAP version you reverted or downgraded to.

If the SP automatic update functionality is disabled (not recommended), after the ONTAP revert or downgrade process is complete, you must manually update the SP firmware to a version that is supported for the ONTAP version you reverted or downgraded to.

Related information

[NetApp BIOS Service Processor Support Matrix](#)

[NetApp Downloads: System Firmware and Diagnostics](#)

Optimal service availability during upgrades

Service availability during ONTAP upgrades can be optimized through planning and configuration. In many cases, upgrades can be completely nondisruptive from a client perspective.

Considerations for services and protocols during upgrades

In general, services based on stateless protocols—such as NFSv3, FC, and iSCSI—are less susceptible to service interruptions during upgrades than session-oriented protocols—such as CIFS and NDMP.

During an upgrade, each node in the cluster must be rebooted (by initiating an HA configuration takeover and giveback) to load the new software. Services based on stateless protocols usually remain available during the nondisruptive upgrade.

Stateless protocols usually include a timeout procedure. For example, if a message is sent and receipt is not acknowledged within a timeout period, a transmission error is assumed to have occurred. In a cluster, if the client's timeout period is greater than the disruption period on the cluster (for example, the amount of time a reboot or HA configuration giveback takes), the client does not perceive a disruption of cluster services.

In session-oriented protocols, there is no concept of timeout to protect the service from disruption. If session-oriented cluster services are disrupted, state information about any operation in progress is lost and the user must restart the operation.

Considerations for stateless protocols

Configurations that include client connections using stateless NAS and SAN protocols generally do not experience adverse effects during upgrades if the clients are configured according to recommended guidelines.

If you are using stateless protocols, consider the following:

- **NFS hard mounts**
No adverse behavior is experienced on the clients during upgrade. Clients might receive some messages similar to the following until the node reboots:
`NFS server not responding, retrying`
In general, read/write directories should be hard-mounted. Hard mounts are the default type of mount.
- **NFS soft mounts**
You should not use soft mounts when there is a possibility of frequent NFS timeouts. Race conditions can occur as a result of these timeouts, which can lead to data corruption. Furthermore, some applications cannot properly handle errors that occur when an NFS operation reaches a timeout using soft mounts.
Situations that can cause frequent timeouts include nondisruptive upgrades or any takeover or giveback event in an HA configuration.
In general, soft mounts should be used only when reading solely from a disk; even then, understand that any soft mount is unreliable.
- **SAN protocols**
No adverse behavior is experienced on FC or iSCSI clients if they are configured according to the recommended guidelines listed in the Interoperability Matrix.

Related information

[*NetApp Interoperability Matrix Tool*](#)

Considerations for session-oriented protocols

Clusters and session-oriented protocols might cause adverse effects on clients and applications in certain areas during upgrades.

If you are using session-oriented protocols, consider the following:

- CIFS
 - Hyper-V and SQL Server over SMB support nondisruptive operations (NDOs). If you configured a Hyper-V or SQL Server over SMB solution, the application servers and the contained virtual machines or databases remain online and provide continuous availability during the ONTAP upgrade.
 - For all other CIFS configurations, client sessions are terminated. You should direct users to end their sessions before you upgrade.
- NFSv4.x
 - NFSv4.x clients will automatically recover from connection losses experienced during the upgrade using normal NFSv4.x recovery procedures. Applications might experience a temporary I/O delay during this process.
- NDMP
 - State is lost and the client user must retry the operation.
- Backups and restores
 - State is lost and the client user must retry the operation.
 - Attention:** Do not initiate a backup or restore during or immediately before an upgrade. Doing so might result in data loss.
- Applications (for example, Oracle or Exchange)
 - Effects depend on the applications. For timeout-based applications, you might be able to change the timeout setting to longer than the ONTAP reboot time to minimize adverse effects.

How firmware is updated during the ONTAP upgrade

Because upgrading ONTAP includes upgrading your firmware, you do not need to update firmware manually. When you perform an ONTAP upgrade, the firmware for your cluster included with the ONTAP upgrade package is copied to each node's boot device, and the new firmware is installed automatically.

Firmware for the following components is updated automatically if the version in your cluster is older than the firmware that is bundled with the ONTAP upgrade package:

- System and diagnostics:
 - BIOS
 - Flash Cache
 - Service Processor (SP)
- Disk
- Disk shelf

If desired, you can also update firmware manually in between ONTAP upgrades.

Related information

[NetApp Downloads: System Firmware and Diagnostics](#)

[NetApp Downloads: Disk Drive and Firmware](#)

[NetApp Downloads: Disk Shelf Firmware](#)

Understanding background disk firmware updates

When a node reboots and there is new disk firmware present, the affected drives are automatically and sequentially taken offline, and the node responds normally to read and write requests.

If any request affects an offline drive, the read requests are satisfied by reconstructing data from other disks in the RAID group, while write requests are written to a log. When the disk firmware update is complete, the drive is brought back online after resynchronizing any write operations that took place while the drive was offline.

During a background disk firmware update, the node functions normally. You see status messages as disks are taken offline to update firmware and brought back online when the firmware update is complete. Background disk firmware updates proceed sequentially for active data disks and for spare disks. Sequential disk firmware updates ensure that there is no data loss through double-disk failure.

Offline drives are marked with the annotation `offline` in the nodeshell `vol status -r` command output. While a spare disk is offline, it cannot be added to a volume or selected as a replacement drive for reconstruction operations. However, a disk would normally remain offline for a very short time (a few minutes at most) and therefore would not interfere with normal cluster operation.

The background disk firmware update is completed unless the following conditions are encountered:

- Degraded aggregates are on the node.
- Disks needing a firmware update are present in an aggregate or plex that is in an offline state.

Automatic background disk firmware updates resume when these conditions are addressed.

Related information

[ONTAP concepts](#)

Copyright

Copyright © 2019 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[*doccomments@netapp.com*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277