



Revert ONTAP

ONTAP 9

NetApp
March 10, 2026

Table of Contents

- Revert ONTAP 1
 - Do I need technical support to revert an ONTAP cluster? 1
 - Supported ONTAP revert paths 1
 - ONTAP revert issues and limitations 2
 - Prepare for an ONTAP revert 3
 - Resources to review before you revert an ONTAP cluster 3
 - System verifications to perform before you revert an ONTAP cluster 4
 - Perform ONTAP version specific pre-revert checks 9
 - Download and install the ONTAP software image 22
 - Download the ONTAP software image 22
 - Install the ONTAP software image 22
 - Revert an ONTAP cluster 24
 - Step 1: Prepare the cluster for reversion 24
 - Step 2: Revert cluster nodes 25
- What to do after an ONTAP revert 32
 - Verify cluster and storage health after an ONTAP revert 32
 - Enable automatic switchover for MetroCluster configurations after an ONTAP revert 35
 - Enable and revert LIFs to home ports after an ONTAP revert 36
 - Enable snapshot policies after an ONTAP revert 38
 - Verify IPv6 firewall entries after an ONTAP revert 39
 - Verify user accounts that can access the Service Processor after reverting to ONTAP 9.8 40

Revert ONTAP

Do I need technical support to revert an ONTAP cluster?

You should contact technical support before you attempt to revert an ONTAP cluster in the following situations:

- A production environment

Do not attempt to revert a production cluster without assistance from technical support.

- You created volumes in ONTAP 9.5 or later and you need to revert to an earlier version.

Volumes using adaptive compression must be uncompressed before reverting.

You can revert new or test clusters without assistance. If you attempt to revert a cluster on your own and experience any of the following issues, you should call technical support:

- The revert fails or cannot finish.
- The revert finishes, but the cluster is unusable in a production environment.
- The revert finishes and the cluster goes into production, but you are not satisfied with its behavior.

Supported ONTAP revert paths

You can directly revert your ONTAP software to only one release earlier than your current ONTAP version. For example, if you are running 9.15.1, you cannot revert directly to 9.13.1. You must first revert to 9.14.1; then perform a separate revert from 9.14.1 to 9.13.1.

Reverting to ONTAP 9.4 or earlier is not supported. You should not revert to unsupported ONTAP versions.

You can use the `system image show` command to determine the version of ONTAP running on each node.

The following supported revert paths refer only to on-premises ONTAP releases. For information about reverting ONTAP in the cloud, see [Reverting or downgrading Cloud Volumes ONTAP](#).



[AFX storage systems](#) do not support reverting ONTAP software.

You can revert from...	To...
ONTAP 9.18.1	ONTAP 9.17.1
ONTAP 9.17.1	ONTAP 9.16.1
ONTAP 9.16.1	ONTAP 9.15.1
ONTAP 9.15.1	ONTAP 9.14.1

You can revert from...	To...
ONTAP 9.14.1	ONTAP 9.13.1
ONTAP 9.13.1	ONTAP 9.12.1
ONTAP 9.12.1	ONTAP 9.11.1
ONTAP 9.11.1	ONTAP 9.10.1
ONTAP 9.10.1	ONTAP 9.9.1
ONTAP 9.9.1	ONTAP 9.8
ONTAP 9.8	ONTAP 9.7
ONTAP 9.7	ONTAP 9.6
ONTAP 9.6	ONTAP 9.5

ONTAP revert issues and limitations

You need to consider the revert issues and limitations before you revert an ONTAP cluster.

- 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.

[AFX storage systems](#) do not support reversion.

- 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.



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.
- In MetroCluster IP systems using switches which are MetroCluster-compliant but not MetroCluster-validated, the reversion from ONTAP 9.7 to 9.6 is disruptive as there is no support for systems using ONTAP 9.6 and earlier.
- Before you revert a node to ONTAP 9.13.1 or earlier, you need to first convert an encrypted SVM root volume to a non-encrypted volume

If you attempt to revert to an ONTAP version that does not support SVM root volume encryption, the system will respond with a warning and block the reversion.

Prepare for an ONTAP revert

Resources to review before you revert an ONTAP cluster

Before you revert an ONTAP cluster, you should confirm hardware support and review resources to understand issues you might encounter or need to resolve.

1. Review the [ONTAP 9 Release Notes](#) for the target release.

The “Important cautions” section describes potential issues that you should be aware of before downgrading or reverting.

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](#)

System verifications to perform before you revert an ONTAP cluster

Before you revert an ONTAP cluster, you should verify your cluster health, storage health, and system time. You should also verify that no jobs are running on the cluster.

Verify cluster health

Before you revert an ONTAP 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
```

In this example, all nodes are healthy and eligible to participate in the cluster.

```
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
```

Enter `y` to continue.

3. 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	<pre>cluster ring show -unitname mgmt</pre>
Volume location database	<pre>cluster ring show -unitname vldb</pre>

To display this RDB process...	Enter this command...
Virtual-Interface manager	<pre>cluster ring show -unitname vifmgr</pre>
SAN management daemon	<pre>cluster ring show -unitname bcomd</pre>

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.
```

4. Return to the admin privilege level:

```
set -privilege admin
```

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

```
event log show -severity informational -message-name scsiblade.*
```

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

```
cluster1::*> event log show -severity informational -message-name
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 ...
```

Related information

Verify storage health

Before you revert an ONTAP cluster, you should verify the status of your disks, aggregates, and volumes.

Steps

1. Verify disk status:

To check for...	Do this...
Broken disks	<p>a. Display any broken disks:</p> <pre>storage disk show -state broken</pre> <p>b. Remove or replace any broken disks.</p>
Disks undergoing maintenance or reconstruction	<p>a. Display any disks in maintenance, pending, or reconstructing states:</p> <pre>storage disk show -state maintenance pending reconstructing</pre> <p>b. Wait for the maintenance or reconstruction operation to finish before proceeding.</p>

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.

```
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.

```
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
```

See the [NetApp Knowledge Base: Volume Showing WAFL Inconsistent](#) on how to address the inconsistent volumes.

Related information

[Disk and aggregate management](#)

Verify the system time

Before you revert an ONTAP cluster, 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
```

```
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.
```

Verify that no jobs are running

Before you revert an ONTAP cluster, you should 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 should 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
```

In this example, there are two jobs queued:

```
cluster1::> job show
```

Job ID	Name	Owning Vserver	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 jobs:

```
job delete -id <job_id>
```

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

```
job show
```

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

```
cluster1::> job show
```

Job ID	Name	Owning Vserver	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

Related information

- [storage disk show](#)

Perform ONTAP version specific pre-revert checks

Pre-revert tasks required for your ONTAP version

Depending upon your ONTAP version, you might need to perform additional preparatory tasks before you begin the revert process.

If you are reverting from ...	Do the following before you start the revert process...
Any ONTAP 9 version	<ul style="list-style-type: none">• Terminate SMB sessions that are not continuously available.• Review reversion requirements for SnapMirror and SnapVault relationships.• Verify deduplicated volumes have enough free space.• Prepare snapshots.• Set the autocommit period for SnapLock volumes to hours.• If you have a Metrocluster configuration, disable automatic unplanned switchover.• Respond to Autonomous Ransomware Protection warnings of abnormal activity before reverting.
ONTAP 9.18.1	<ul style="list-style-type: none">• If automatic enablement has been set for ARP as part of an ONTAP 9.18.1 upgrade, you'll need to disable it.
ONTAP 9.17.1	<ul style="list-style-type: none">• If you have enabled the ONTAP ARP feature for SAN, disable it.
ONTAP 9.16.1	<ul style="list-style-type: none">• If you have TLS configured for NVMe/TCP connections, disable the TLS configuration on the NVME hosts.• If you have extended qtree performance monitoring enabled, disable it.• If you are using CORS to access your ONTAP s3 buckets, remove the CORS configuration.
ONTAP 9.14.1	<p>If you have enabled trunking for client connections, disable trunking on any NFSv4.1 servers.</p>

If you are reverting from ...	Do the following before you start the revert process...
ONTAP 9.12.1	<ul style="list-style-type: none"> • If you have configured S3 client access for NAS data, remove the S3 NAS bucket configuration. • If you are running the NVMe protocol and have configured in-band authentication, disable in-band authentication. • If you have a Metrocluster configuration, disable IPsec.
ONTAP 9.11.1	If you have configured Autonomous Ransomware Protection (ARP), check the ARP licensing .
ONTAP 9.6	If you have SnapMirror synchronous relationships, prepare the relationships for revert .

Any ONTAP 9 version

Terminate certain SMB sessions before reverting ONTAP

Before you revert an ONTAP cluster from any version of ONTAP 9, you should identify and gracefully terminate any SMB sessions that are not continuously available.

Continuously available SMB 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 SMB sessions that are not continuously available:

```
vserver cifs session show -continuously-available No -instance
```

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

```
cluster1::> vserver cifs session show -continuously-available No
-instance
```

```
Node: node1
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: CIFSLAB\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 SMB session that you identified:

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

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

```
Node:      node1
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.
```

ONTAP revert 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 revert 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 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.



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.

Learn more about `snapmirror restore` in the [ONTAP command reference](#).

Verify free space for deduplicated volumes before reverting ONTAP

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

The volume must have enough space to accommodate the savings that were achieved through the inline detection of blocks of zeros. See the [NetApp Knowledge Base: How to see space savings from deduplication, compression, and compaction in ONTAP 9](#).

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. View the progress of the efficiency operations that are running on the volumes:

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

2. Stop all active and queued deduplication operations:

```
volume efficiency stop -vserver <svm_name> -volume <volume_name> -all
```

3. Set the privilege level to advanced:

```
set -privilege advanced
```

4. Downgrade the efficiency metadata of a volume to the target version of ONTAP:

```
volume efficiency revert-to -vserver <svm_name> -volume <volume_name> -version <version>
```

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

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



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. Monitor the progress of the downgrade:

```
volume efficiency show -vserver <svm_name> -op-status Downgrading
```

6. If the revert does not succeed, display the instance to see why the revert failed.

```
volume efficiency show -vserver <svm_name> -volume <volume_name> -instance
```

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

```
set -privilege admin
```

Learn more about [Logical storage management](#).

Prepare Snapshots before reverting an ONTAP cluster

Before you revert an ONTAP cluster from any version of ONTAP 9, you must disable all snapshot policies and delete any Snapshots that were created after upgrading to the current release.

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 policies for all data SVMs:

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

2. Disable snapshot policies for each node's aggregates:

- a. Identify the node's aggregates:

```
run -node <nodename> -command aggr status
```

- b. Disable the snapshot policy for each aggregate:

```
run -node <nodename> -command aggr options aggr_name nosnap on
```

- c. Repeat this step for each remaining node.

3. Disable snapshot policies for each node's root volume:

- a. Identify the node's root volume:

```
run -node <node_name> -command vol status
```

You identify the root volume by the word **root** in the **Options** column of the `vol status` command output.

```
vs1::> run -node node1 vol status
```

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

- b. Disable the snapshot policy on the root volume:

```
run -node <node_name> vol options root_volume_name nosnap on
```

c. Repeat this step for each remaining node.

4. Delete all snapshots 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 snapshots:

```
volume snapshot prepare-for-revert -node <node_name>
```

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

If any snapshots cannot be deleted, the command fails and notifies you of any required actions you must take before the snapshots 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.

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

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

d. Verify that the snapshots have been deleted:

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

e. If any newer-version snapshots remain, force them to be deleted:

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

f. Repeat these steps for each remaining node.

g. Return to the admin privilege level:

```
set -privilege admin
```



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

Set autocommit periods for SnapLock volumes before reverting ONTAP

Before you revert an ONTAP cluster from any version of ONTAP 9, the value of the autocommit period for SnapLock volumes must be set in hours, not days. You should 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 <vserver_name> -volume <volume_name>  
-autocommit-period value hours
```

Disable automatic unplanned switchover before reverting MetroCluster configurations

Before reverting a MetroCluster configuration running any version of ONTAP 9, 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](#)

Resolve activity warnings in Autonomous Ransomware Protection (ARP) before an ONTAP revert

Before you revert to ONTAP 9.17.1 or earlier, you should respond to any abnormal activity warnings reported by Autonomous Ransomware Protection (ARP) and delete any associated ARP screenshots.

Before you begin

You need "Advanced" privileges to delete ARP snapshots.

Steps

1. Respond to any abnormal activity warnings reported by [ARP](#) and resolve any potential issues.
2. Confirm the resolution of these issues before reverting by selecting **Update and Clear Suspect File Types** to record your decision and resume normal ARP monitoring.
3. List any ARP screenshots associated with the warnings by running the following command:

```
volume snapshot snapshot show -fs-version 9.18
```

4. Delete any ARP screenshots associated with the warnings:



This command deletes all snapshots that have the format used by the current version of ONTAP, potentially not just ARP snapshots. Ensure that you have taken any necessary actions for all snapshots that will be removed before running this command.

```
volume snapshot prepare-for-revert -node <node_name>
```

ONTAP 9.18.1

Disable automatic enablement of Autonomous Ransomware Protection before reverting from ONTAP 9.18.1

If you upgraded volumes to ONTAP 9.18.1, ONTAP ARP automatic enablement might have been set for your volumes after a brief grace period (12 hours). It's recommended that you disable this automatic enablement setting on volumes upgraded to ONTAP 9.18.1 before reverting to ONTAP 9.17.1 or earlier.

Steps

1. Determine if the automatic enablement option has been activated on volumes that have been upgraded to ONTAP 9.18.1 or later:

```
security anti-ransomware auto-enable show
```

2. Disable the automatic enablement option for ransomware protection on all volumes on the SVM:

```
security anti-ransomware volume disable -volume * -auto-enabled-volumes  
-only true
```

ONTAP 9.17.1

Disable Autonomous Ransomware Protection on SAN volumes before reverting from ONTAP 9.17.1

The ONTAP ARP feature for SAN volumes is not supported in ONTAP 9.16.1 and earlier.

It's recommended that you disable ARP on SAN volumes before reverting to ONTAP 9.16.1 or earlier to prevent the feature from staying active and using CPU and disk resources without performing any actual detection on the reverted version.

Example 1. Steps

System Manager

1. Select **Storage > Volumes**, then select the name of the volume.
2. In the **Security** tab of the **Volumes** overview, select **Status** to switch from Enabled to Disabled.

CLI

1. Disable ransomware protection on a volume:

```
security anti-ransomware volume disable -volume <vol_name> -vserver  
<svm_name>
```

ONTAP 9.16.1

Disable TLS on NVMe hosts before reverting from ONTAP 9.16.1

If you have TLS secure channel for NVMe/TCP connections configured on an NVMe host, you need to disable it before you revert your cluster from ONTAP 9.16.1.

Steps

1. Remove the TLS secure channel configuration from the host:

```
vserver nvme subsystem host unconfigure-tls-for-revert -vserver  
<svm_name> -subsystem <subsystem> -host-nqn <host_nqn>
```

This command removes the host from the subsystem, and then recreates the host in the subsystem without the TLS configuration.

2. Verify that TLS secure channel is removed from the host:

```
vserver nvme subsystem host show
```

Disable extended Qtree performance monitoring before reverting from ONTAP 9.16.1

Beginning with ONTAP 9.16.1, you can use the ONTAP REST API to access the extended qtree monitoring capabilities which includes latency metrics and historical statistics. If extended qtree monitoring is enabled on any qtrees, before you revert from 9.16.1, you must set `ext_performance_monitoring.enabled` to `false`.

You work with the `ext_performance_monitoring.enabled` setting by using the `/api/storage/qtrees`

endpoint. Use GET to retrieve the current value, POST to set it when creating a new qtree, and PATCH to change it on an existing qtree.

Learn more about [reverting clusters with extended qtree performance monitoring](#).

Remove CORS configuration before reverting from ONTAP 9.16.1

If you are using Cross-Origin Resource Sharing (CORS) to access ONTAP S3 buckets, you must remove it before you revert from ONTAP 9.16.1.

Learn more about [reverting ONTAP clusters with using CORS](#).

ONTAP 9.14.1

Disable NFSv4.1 session trunking before reverting from ONTAP 9.14.1

If you have enabled trunking for client connections, you must disable trunking on any NFSv4.1 servers before reverting from ONTAP 9.14.1.

When you enter the `revert-to` command, you will see a warning message advising you to disable trunking before proceeding.

After reverting to an ONTAP 9.13.1, the clients using trunked connections fall back to using a single connection. Their data throughput will be affected, but there will be no disruption. The revert behavior is the same as modifying the NFSv4.1 trunking option for the SVM from enabled to disabled.

Steps

1. Disable trunking on the NFSv4.1 server:

```
vserver nfs modify -vserver _svm_name_ -v4.1-trunking disabled
```

2. Verify that NFS is configured as desired:

```
vserver nfs show -vserver _svm_name_
```

ONTAP 9.12.1

Remove S3 NAS bucket configuration before reverting from ONTAP 9.12.1

If you have configured S3 client access for NAS data, you should use the ONTAP command line interface (CLI) to remove the NAS bucket configuration and to remove any name mappings (S3 users to Windows or Unix users) before reverting from ONTAP 9.12.1.

About this task

The following tasks are completed in the background during the revert process.

- Remove all partially completed singleton object creations (that is, all entries in hidden directories).

- Remove all hidden directories; there might be one on for each volume that is accessible from the root of the export mapped from the S3 NAS bucket.
- Remove the upload table.
- Delete any default-unix-user and default-windows-user values for all configured S3 servers.

Steps

1. Remove S3 NAS bucket configuration:

```
vserver object-store-server bucket delete -vserver <svm_name> -bucket <s3_nas_bucket_name>
```

Learn more about `vserver object-store-server bucket delete` in the [ONTAP command reference](#).

2. Remove name mappings for UNIX:

```
vserver name-mapping delete -vserver <svm_name> -direction s3-unix
```

Learn more about `vserver name-mapping delete` in the [ONTAP command reference](#).

3. Remove name mappings for Windows:

```
vserver name-mapping delete -vserver <svm_name> -direction s3-win
```

4. Remove the S3 protocols from the SVM:

```
vserver remove-protocols -vserver <svm_name> -protocols s3
```

Learn more about `vserver remove-protocols` in the [ONTAP command reference](#).

Disable NVMe in-band authentication before reverting from ONTAP 9.12.1

If you are running the NVME protocol, you must disable in-band authentication before you revert your cluster from ONTAP 9.12.1. If in-band authentication using DH-HMAC-CHAP is not disabled, revert will fail.

Steps

1. Remove the host from the subsystem to disable DH-HMAC-CHAP authentication:

```
vserver nvme subsystem host remove -vserver <svm_name> -subsystem <subsystem> -host-nqn <host_nqn>
```

2. Verify that the DH-HMAC-CHAP authentication protocol is removed from the host:

```
vserver nvme subsystem host show
```

3. Add the host back to the subsystem without authentication:

```
vserver nvme subsystem host add vserver <svm_name> -subsystem  
<subsystem> -host-nqn <host_nqn>
```

Disable IPsec in MetroCluster configurations before reverting from ONTAP 9.12.1

Before reverting a MetroCluster configuration from ONTAP 9.12.1, you must disable IPsec.

A check is performed before revert to ensure there are no IPsec configurations within the MetroCluster configuration. You must remove any IPsec configurations present and disable IPsec before continuing with the revert. Reverting ONTAP is blocked if IPsec is enabled, even when you have not configured any user policies.

ONTAP 9.11.1

Check Autonomous Ransomware Protection licensing before reverting from ONTAP 9.11.1

If you have configured Autonomous Ransomware Protection (ARP) and you revert from ONTAP 9.11.1 to ONTAP 9.10.1, you might experience warning messages and limited ARP functionality.

In ONTAP 9.11.1, the Anti-ransomware license replaced the Multi-Tenant Key Management (MTKM) license. If your system has the Anti_ransomware license but no MT_EK_MGMT license, you will see a warning during revert that ARP cannot be enabled on new volumes upon revert.

The volumes with existing protection will continue to work normally after revert, and ARP status can be displayed using the ONTAP CLI. System Manager cannot show ARP status without the MTKM license.

Therefore, if you want ARP to continue after reverting to ONTAP 9.10.1, be sure the MTKM license is installed before reverting. [Learn about ARP licensing.](#)

ONTAP 9.6

Considerations for reverting systems from ONTAP 9.6 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.

ONTAP 9.5 does not support NFSv4 and SMB.

- 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 snapshots 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 snapshots from ONTAP 9.5 are automatically replaced by the common snapshots in ONTAP 9.6. Therefore, you cannot resynchronize the SnapMirror synchronous relationship after reverting if the common snapshots from ONTAP 9.5 are not available.

Download and install the ONTAP software image

Before you revert your current ONTAP software, you must download the target software version from the NetApp Support site and then install it.

Download the ONTAP software image

Software images are specific to platform models. You must obtain the correct image for your cluster. Software images, firmware version information, and the latest firmware for your platform model are available on the NetApp Support Site. Software images include the latest version of system firmware that was available when a given version of ONTAP was released.



If you are reverting a system with NetApp Volume Encryption from 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 revert 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 example, `97_q_image.tgz`) from the NetApp Support Site

You can copy the image to the directory on the HTTP server or FTP server from which the image will be served or to a local folder.

Install the ONTAP software image

After downloading the target ONTAP software image from the NetApp support site, install it on the cluster nodes.

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

The advanced prompt (`*>`) appears.

2. Enter `y` to continue when prompted .

3. Install the software image:

- For standard configurations or a two-node MetroCluster configuration enter the following command:

```
system node image update -node * -package  
<http://example.com/downloads/image.tgz> -replace-package true  
-replace {image1|image2} -background true -setdefault 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. This command also 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.

- For a four-node or eight-node MetroCluster configuration, enter the following command on both clusters:

```
system node image update -node * -package  
<http://example.com/downloads/image.tgz> -replace-package true  
-replace {image1|image2} -background true -setdefault false
```

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. This command also uses an extended query to change the target software image, which is installed as the alternate image on each node.

4. Enter `y` to continue when prompted.

5. 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.

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.
```

Related information

- [system node image update](#)

Revert an ONTAP cluster

Reverting an ONTAP cluster is disruptive. You must take the cluster offline for the duration of the reversion. You should not revert a production cluster without assistance from technical support.

To revert a new or test cluster, you must disable storage failover and data LIFs and address reversion preconditions; then you must revert the cluster and file system configuration on each node in the cluster.

Before you begin.

- You should have completed the [pre-revert verifications](#).
- You should have completed the required [pre-checks for your specific ONTAP version](#).
- You should have [downloaded and installed the target ONTAP software image](#).

Step 1: Prepare the cluster for reversion

Before you revert any of your cluster nodes, you should verify that your target ONTAP image is installed and you should disable all the data LIFs in the cluster.

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

Enter **y** when prompted to continue.

2. Verify that the target ONTAP software is installed:

```
system image show
```

The following example shows that version 9.13.1 is installed as the alternate image on both nodes:

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

Node	Image	Is Default	Is Current	Version	Install Date
node0					
	image1	true	true	9.14.1	MM/DD/YYYY TIME
	image2	false	false	9.13.1	MM/DD/YYYY TIME
node1					
	image1	true	true	9.14.1	MM/DD/YYYY TIME
	image2	false	false	9.13.1	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. Determine if you have inter-cluster flexcache relationships:

```
flexcache origin show-caches -relationship-type inter-cluster
```

5. If inter-cluster flexcaches are present, disable the data lifs on the cache cluster:

```
network interface modify -vserver <vserver_name> -lif <lif_name> -status  
-admin down
```

Step 2: Revert cluster nodes

To revert your cluster, you need to revert the first node in an HA pair, then revert the partner node. You then repeat this process for each HA pair in your cluster until all nodes are reverted. If you have a MetroCluster configuration, you need to repeat these steps for both the clusters in the configuration.

4 or more nodes

Steps

1. 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.

2. Disable storage failover for the nodes in the HA pair:

```
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.

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

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

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

```
system image show
```

The following example shows that version 9.13.1 is set as the default image on node0:

```
cluster1::*> system image show

      Is      Is      Install
Node  Image   Default Current Version  Date
-----
node0
      image1  false   true   9.14.1  MM/DD/YYYY TIME
      image2  true    false  9.13.1  MM/DD/YYYY TIME
node1
      image1  true    true   9.14.1  MM/DD/YYYY TIME
      image2  false   false  9.13.1  MM/DD/YYYY TIME
4 entries were displayed.
```

5. 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 disabling the snapshot policy or deleting snapshots that were created after upgrading to the later version of ONTAP.

The `-version` option refers to the ONTAP release to which you are reverting. For example, if you are reverting from 9.14.1 to 9.13.1, the correct value of the `-version` option is 9.13.1.

6. Revert the cluster configuration of the node:

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

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

7. Wait for the login prompt; then enter **No** when you are asked if you want to login to the systemshell.

It could take up to 30 minutes or more for the login prompt to appear.

8. Log in to the clustershell with admin.

9. 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.

10. 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.



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

If `AUTOBOOT` is true, when the command finishes, the node will reboot to ONTAP.

If `AUTOBOOT` is false, when the command finishes, the `LOADER` prompt is displayed. Enter `yes` to revert; then use `boot_ontap` to manually reboot the node.

11. After the node has rebooted, confirm that the new software is running:

```
system node image show
```

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

```

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

```

- Verify that the revert status for the node is complete:

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

The status should be listed as "complete", "not needed", or "there are no table entries returned."

- Repeat these steps on the other node in the HA pair; then repeat these steps for each additional HA pair.

If you have a MetroCluster Configuration, you need to repeat these steps on both clusters in the configuration

- After all nodes have been reverted, reenable high availability for the cluster:

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

2-node cluster

- 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.

- Disable cluster high availability (HA):

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

- Disable storage failover:

```
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.

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

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

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

```
system image show
```

The following example shows that version 9.13.1 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	9.14.1	MM/DD/YYYY TIME
	image2	true	false	9.13.1	MM/DD/YYYY TIME
node1	image1	true	true	9.14.1	MM/DD/YYYY TIME
	image2	false	false	9.13.1	MM/DD/YYYY TIME

4 entries were displayed.

6. Check whether the node currently holds epsilon:

```
cluster show -node <nodename>
```

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

- a. 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 <nodename> -epsilon false
```

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

```
cluster modify -node <node_partner_name> -epsilon true
```

7. 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 conditions that must be addressed before reverting, such as disabling the snapshot policy or deleting snapshots that were created after upgrading to the later version of ONTAP.

The `-version` option refers to the ONTAP release to which you are reverting. Only the first two values of the ONTAP version are required. For example, if you are reverting from 9.14.1 to 9.13.1, the correct value of the `-version` option is 9.13.

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

8. Revert the cluster configuration of the node:

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

9. Wait for the login prompt; then enter `N` when you are asked if you want to login to the systemshell.

It could take up to 30 minutes or more for the login prompt to appear.

10. Log in to the clustershell with admin.

11. 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.

12. 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.



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

If AUTOBOOT is true, when the command finishes, the node will reboot to ONTAP.

If AUTOBOOT is false, when the command finishes the LOADER prompt is displayed. Enter `yes` to revert; then use `boot_ontap` to manually reboot the node.

13. After the node has rebooted, confirm that the new software is running:

```
system node image show
```

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

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

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

4 entries were displayed.

14. Verify that the revert status is complete for the node:

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

The status should be listed as "complete", "not needed", or "there are no table entries returned."

15. Repeat these steps on the other node in the HA pair.
16. After both nodes have been reverted, reenabling high availability for the cluster:

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

17. Reenable storage failover on both nodes:

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

Related information

- [storage failover modify](#)

What to do after an ONTAP revert

Verify cluster and storage health after an ONTAP revert

After you revert an ONTAP cluster, you should verify that the nodes are healthy and eligible to participate in the cluster, and that the cluster is in quorum. You should also verify the status of your disks, aggregates, and volumes.

Verify cluster health

Steps

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

```
cluster show
```

In this example, the cluster is healthy and all nodes are eligible to participate in the cluster.

```
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
```

Enter `y` to continue.

3. 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>

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.
```

4. Return to the admin privilege level:

```
set -privilege admin
```

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

```
event log show -severity informational -message-name scsiblade.*
```

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

```
cluster1::*> event log show -severity informational -message-name
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 ...

Related information

[System administration](#)

Verify storage health

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

Steps

1. Verify disk status:

To check for...	Do this...
Broken disks	<ol style="list-style-type: none"> a. Display any broken disks: <pre>storage disk show -state broken</pre> b. Remove or replace any broken disks.
Disks undergoing maintenance or reconstruction	<ol style="list-style-type: none"> a. Display any disks in maintenance, pending, or reconstructing states: <pre>storage disk show -state maintenance pending reconstructing</pre> 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.

```
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.

```
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
```

See the [NetApp Knowledge Base: Volume Showing WAFL Inconsistent](#) on how to address the inconsistent volumes.

Verify client access (SMB and NFS)

For the configured protocols, test access from SMB and NFS clients to verify that the cluster is accessible.

Related information

- [Disk and aggregate management](#)
- [storage disk show](#)

Enable automatic switchover for MetroCluster configurations after an ONTAP revert

After reverting an ONTAP MetroCluster configuration, you must enable automatic unplanned switchover to ensure that the MetroCluster configuration is fully operational.

Steps

1. Enable automatic unplanned switchover:

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

2. Validate the MetroCluster configuration:

```
metrocluster check run
```

Enable and revert LIFs to home ports after an ONTAP revert

During a reboot, some LIFs might have been migrated to their assigned failover ports. After you revert an ONTAP cluster, you must enable and revert any LIFs that are not on their home ports.

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
```

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.

2. Enable the data LIFs:

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

3. Revert LIFs to their home ports:

```
network interface revert *
```

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

```
network interface show
```

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

```

cluster1::> network interface show -vserver vs0
      Logical      Status      Network      Current
Current Is
Vserver      Interface  Admin/Oper  Address/Mask  Node      Port
Home
-----
vs0
      data001      up/up      192.0.2.120/24  node0      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.

```

Related information

- [network interface](#)

Enable snapshot policies after an ONTAP revert

After reverting to an earlier version of ONTAP, you must enable snapshot policies to start creating snapshots again.

You are reenabling the snapshot schedules that you disabled before you reverted to an earlier version of ONTAP.

Steps

1. Enable snapshot policies for all data SVMs:

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

```
snapshot policy modify pg-rpo-hourly -enable true
```

2. For each node, enable the snapshot policy of the root volume:

```
run -node <node_name> vol options <volume_name> nosnap off
```

Verify IPv6 firewall entries after an ONTAP revert

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
```

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 -policy <policy_name> -service  
ssh -action allow -ip-list <ip_list>
```

3. Apply the new policy to the LIF to allow access to a network service:

```
network interface modify -vserve <svm_name> -lif <lif_name> -firewall  
-policy <policy_name>
```

Verify user accounts that can access the Service Processor after reverting to ONTAP 9.8

In ONTAP 9.9.1 and later the the `-role` parameter for user accounts is changed to `admin`. If you created user accounts on ONTAP 9.8 or earlier, upgraded to ONTAP 9.9.1 or later and then reverted back to ONTAP 9.8, the `-role` parameter is restored to its original value. You should verify that the modified values are acceptable.

During revert, if the role for an SP user has been deleted, the "rbac.spuser.role.notfound" EMS message will be logged.

For more information, see [Accounts that can access the SP](#).

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.