



Preparing for copy-based transition

ONTAP 7-Mode Transition

NetApp
February 11, 2024

This PDF was generated from https://docs.netapp.com/us-en/ontap-7mode-transition/copy-based/concept_requirements_for_copy_based_transition.html on February 11, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Preparing for copy-based transition 1
 - Requirements for copy-based transition 1
 - Port requirements for communicating with the 7-Mode Transition Tool 2
 - Restrictions for transition 4
 - Preparing the 7-Mode system for transition 5
 - Preparing the network for transition 6
 - Preparing the cluster for transition 7
 - Preparing 7-Mode aggregates and volumes for transition 10
 - Support for transitioning SnapLock volumes 16
 - Preparing to transition name services 18
 - Preparing for NFS transition. 21
 - Preparing for SMB/CIFS transition. 28
 - Preparing for MetroCluster configuration transition 35
 - Preparing for SAN transition 36
 - Preparing data protection features for transition 40
 - Guidelines for deciding when to perform cutover. 42
 - Impact of takeover and giveback on transition. 43

Preparing for copy-based transition

Before initiating a data copy operation from 7-Mode to ONTAP, you must understand the requirements and restrictions for migration, and complete certain tasks on the 7-Mode system and the cluster.

You must ensure that the following requirements are met before transition:

- The 7-Mode and ONTAP systems must be reachable from the host on which the tool is installed.
- The 7-Mode systems must be running the supported Data ONTAP versions.
- SnapMirror must be licensed on the 7-Mode system.
- Required feature licenses, if they exist on the 7-Mode system, must be installed on the cluster.
- The NTP server must be configured and the time must be synchronized across the 7-Mode system and cluster.
- All preparatory tasks on the 7-Mode system must be completed.
- All preparatory tasks on the cluster must be completed.

Related information

[Transition preparation checklist](#)

[7MTT v2.0/Transitioned Data ONTAP features](#)

[NetApp Interoperability Matrix Tool](#)

Requirements for copy-based transition

You must be aware of the ONTAP release requirements, licensing requirements, and 7-Mode Transition Tool requirements for copy-based transition.


• Data ONTAP 7-Mode source systems

For a list of the 7-Mode releases supported for migration by the 7-Mode Transition Tool, see the [NetApp Interoperability Matrix Tool](#).

• ONTAP target systems

Copy-based transitions are supported to these ONTAP target releases.

If your transition target is running ...	You must use this 7-Mode Transition Tool version ...
ONTAP 9.10.1, ONTAP 9.11.1 or earlier supported release	3.5.0
ONTAP 9.9.1 or earlier supported release	3.4.0
ONTAP 9.8 or earlier supported release	3.3.3

If your transition target is running ...	You must use this 7-Mode Transition Tool version ...
ONTAP 9.7P2 or later 9.7 P release  Earlier 9.7 releases are not supported.	3.3.2
ONTAP 9.6P7 or later 9.6 P release  Earlier 9.6 releases are not supported.	3.3.2
ONTAP 9.5 or earlier ONTAP 9 release	3.3.2 or 3.3.1
Clustered Data ONTAP 8.1.4P4 or later 8.x release	3.3.2 or 3.3.1

- **Licensing requirements**

SnapMirror must be licensed on the 7-Mode storage system. If the 7-Mode system does not have a SnapMirror license, you can obtain a temporary SnapMirror license for transition from your sales representative.

SnapLock must be licensed on the destination cluster if Chain of Custody verification must be performed.

- **7-Mode Transition Tool service**

For the data copy schedules to take effect, the 7-Mode Transition Tool service must be always up and running on the Windows or Linux system on which the tool is installed. However, it does not require the web interface to be active or open for the schedules to take effect. You can close the web interface and re-log in whenever required.

- **Storage, host, and FC switch version requirements for transition assessment**

For the list of 7-Mode versions, hosts, and FC switches that are supported for assessment by the 7-Mode Transition Tool, see the [NetApp Interoperability Matrix Tool](#).

Port requirements for communicating with the 7-Mode Transition Tool

The 7-Mode Transition Tool communicates with the 7-Mode system and the cluster over certain ports. You must ensure that these ports on the 7-Mode system and the cluster are open to allow communication with the 7-Mode Transition Tool.

Ports that must be open on the 7-Mode systems

The 7-Mode Transition Tool communicates with the 7-Mode systems by using HTTPS on port 443.

The following ports are required by the cluster to communicate with the 7-Mode systems for SnapMirror

replication:

- 10565/TCP
- 10566/TCP
- 10567/TCP
- 10568/TCP
- 10569/TCP
- 10670/TCP

Ports that must be open on the cluster

The 7-Mode Transition Tool communicates with the cluster by using HTTPS on port 443.

The following ports are required by the 7-Mode systems to communicate with the cluster for SnapMirror replication:

- 10565/TCP
- 10566/TCP
- 10567/TCP
- 10568/TCP
- 10569/TCP
- 10670/TCP
- 11105/TCP

Additionally, the 7-Mode Transition Tool performs a ping operation from the intercluster LIFs to the data copy IP address of the 7-Mode system to verify reachability.

Ports that must be open on the 7-Mode Transition Tool

Port 8444 of the 7-Mode Transition Tool must be open for the web interface.

To transition netgroups and CIFS local users and groups, the following requirements must be met:

- Port 8088 of the 7-Mode Transition Tool must be available.

For an alternative to port 8088, you must change the port specified by the `tool.http.port` parameter in the `transition-tool.conf` file of the 7-Mode Transition Tool installation directory.



You must restart the 7-Mode Transition Tool service after changing the port in the configuration file.

- Each node in the cluster must have at least one data LIF configured for the target SVM.
- All SVM data LIFs must be able to communicate with the 7-Mode Transition Tool port 8088 or the port specified by the `tool.http.port` parameter in the `transition-tool.conf` file.



You must verify that firewalls do not block this traffic.

Restrictions for transition

You must be aware of certain restrictions for transitioning some 7-Mode volumes and configurations.

- No volume within the same project can cut over until all volumes in the same project have completed their baseline transfers.
- If you want to transition 7-Mode primary and secondary volumes when both the 7-Mode source and destination are running Data ONTAP 7.3.x or 8.0.x, you must start transitioning the 7-Mode secondary volume only when there are no data updates from the 7-Mode primary to the 7-Mode secondary volume.

You must verify that the data update schedules for the 7-Mode primary volume to the 7-Mode secondary volume do not conflict with the schedules for the 7-Mode secondary volume to the ONTAP secondary volume.

- You must not initiate a transition while the aggregates on either the 7-Mode system or cluster are upgrading from 32-bit to 64-bit format; otherwise the transition fails.
- The 7-Mode Transition tool does not transition a volume with a qtree that is the destination of a qtree SnapMirror relationship.

The qtree SnapMirror relationship must be broken before the volume can be transitioned.

- You cannot transition a fanout SnapMirror relationship (a primary volume that is in SnapMirror relationships with more than one secondary volume in different controllers) by using the 7-Mode Transition Tool web interface.

To transition the SnapMirror relationships in a fanout configuration, you must use the 7-Mode Transition Tool CLI. You should create separate projects for each secondary volume, complete the transition of the secondary projects, and then create and complete the transition of the primary volume.

- You cannot transition volumes from different vFiler units or from different 7-Mode controllers to the same SVM at the same time.

You must complete the transition of volumes from a given vFiler unit or 7-Mode controller before you can start the transition of volumes from another vFiler unit or 7-Mode controller.

- The 7-Mode Transition tool does not transition a vFiler unit as a single entity.

However, you can transition all of the volumes in a vFiler unit by selecting them as a part of one or more projects.

- The 7-Mode Transition tool does not transition the root volume of a vFiler unit if the root volume is based on a qtree that belongs to the default vFiler unit.
- The 7-Mode Transition tool does not transition a volume with a qtree if the volume and qtree are owned by different vFiler units.

Transitioning such a volume causes the qtree to become inaccessible.

The precheck operation displays information about some of these restrictions.

Preparing the 7-Mode system for transition

Before starting a transition, you must complete certain tasks on the 7-Mode system, such as adding the SnapMirror license, enabling the 7-Mode system to communicate with the target cluster, and enabling TLS.

All the 7-Mode volumes that you want to transition must be online.

Steps

1. Add and enable the SnapMirror license on the 7-Mode system:

- a. Add the SnapMirror license on the 7-Mode system:

```
license add license_code
```

license_code is the license code you purchased.

- b. Enable the SnapMirror functionality: **+ options snapmirror.enable on**

2. Configure the 7-Mode system and the target cluster to communicate with each other by choosing one of the following options:
 - Set the `snapmirror.access` option to `all`.
 - Set the value of the `snapmirror.access` option to the IP addresses of all the intercluster LIFs on the cluster.
 - If the `snapmirror.access` option is `legacy` and the `snapmirror.checkip.enable` option is `off`, add the SVM name to the `/etc/snapmirror.allow` file.
 - If the `snapmirror.access` option is `legacy` and the `snapmirror.checkip.enable` option is `on`, add the IP addresses of the intercluster LIFs to the `/etc/snapmirror.allow` file.
3. If HTTPS is not enabled on the storage system, enable HTTPS:

```
options httpd.admin.ssl.enable on
```

HTTPS is enabled by default.

4. Enable TLS on the 7-Mode storage systems for enabling the 7-Mode Transition Tool to communicate with the 7-Mode systems:
 - a. If SSL is not already enabled on the storage system, set up and start SSL:

```
secureadmin setup ssl
```

SSL is set up for the storage systems by default. If SSL has been previously set up for the storage system, you are asked whether you want to continue. You can exit the SSL setup if you do not want to make any changes.

- b. Enable SSL:

```
options ssl.enable on
```

This option must be enabled for allowing communication over TLS.

c. Enable TLS:

```
options tls.enable on
```

d. Disable SSLv2 and SSLv3 on the 7-Mode system:

```
options ssl.v2.enable off
```

```
options ssl.v3.enable off
```

The 7-Mode Transition Tool uses TLS or SSL protocols for communicating with the 7-Mode storage systems. The tool communicates with the storage system using the TLS protocol if TLS is enabled on the storage system. If TLS is disabled and SSLv3 is enabled on a storage system, the tool uses SSLv3 to communicate with the storage system.

+ IMPORTANT: The best practice is to enable TLS and disable SSLv2 and SSLv3 in order to avoid security vulnerabilities.

5. Depending on the Data ONTAP version of your 7-Mode system, perform the following steps:

a. Allow SnapMirror traffic on all the interfaces:

```
options interface.blocked.snapmirror ""
```

b. If you are running Data ONTAP version 7.3.7, 8.0.3, or 8.1 and you are using the IP address of the e0M interface as the management IP address to interact with 7-Mode Transition Tool, allow data traffic on the e0M interface:

```
options interface.blocked.mgmt_data_traffic off
```

6. If you have set the I2P, read allocations, or NVFAIL options on the volume, perform the following steps:

a. Verify that other operations are not impacted if these options are disabled.

b. Disable the options:

```
vol options vol_name no_i2p off
```

```
vol options vol_name read_realloc off
```

```
vol options vol_name nvfail off
```

Preparing the network for transition

You must prepare the data network of the cluster for transition by creating logical ports (VLANs and interface groups).

The NTP server must be configured and the time must be synchronized across the 7-Mode systems and cluster.

Steps

1. Create VLANs or interface groups on the target cluster nodes, if required:

```
network port vlan create
```


or

network port ifgrp create

To provide network connectivity after transition, you should transition the 7-Mode IP addresses to a similar network topology in ONTAP. For example, if the 7-Mode IP addresses are configured on physical ports, the IP addresses should be transitioned to appropriate physical ports in ONTAP. Similarly, IP addresses configured on VLAN ports or interface groups should be transitioned to appropriate VLAN ports or interface groups in ONTAP.

2. If you want SVMs in the non-default IPspace, create the required IPspaces:

network ipspace create

The 7-Mode IP addresses or the new LIFs that are selected for transition are created in the IPspace of the mapped SVM.



IPv6 addresses cannot be transitioned and must be configured manually post-transition.

Related information

[Network and LIF management](#)

Considerations for transitioning 7-Mode IP addresses

You must be aware of certain considerations when transitioning 7-Mode IP addresses to storage virtual machines (SVMs) in ONTAP.

- You can transition existing 7-Mode IP addresses or specify new IP addresses to be configured on the SVM by using the 7-Mode Transition Tool.
 - Existing 7-Mode IP addresses are created on the SVM in the administrative `down` state in the apply configuration (precutover) phase.
 - New IP addresses are created on the SVM in the administrative `up` state in the apply configuration (precutover) phase.
- IPv6 addresses cannot be transitioned and must be manually configured after the transition.
- iSCSI and FC LIFs are not transitioned and must be manually configured after the transition.

Preparing the cluster for transition

Before transition, you must ensure that the cluster meets requirements such as allowing HTTPS, setting up intercluster LIFs, and verifying the network connectivity for transition.

- The cluster and the SVM must already be set up.

[Software setup](#)

The target SVM must not be in an SVM disaster recovery relationship.

- The cluster must be healthy and none of the nodes must be in takeover mode.

- The target aggregates that will contain the transitioned volumes must have an SFO policy.
- The aggregates must be on nodes that have not reached the maximum volume limit.
- If you want to transition volumes from a 32-bit aggregate of a 7-Mode system to a 64-bit aggregate of a Data ONTAP 8.2.x cluster, you must have provided an additional 5 percent space in the destination aggregate.

The additional space is required to upgrade the transitioned volume to 64-bit format.

Disk and aggregate management

- For establishing an SVM peer relationship when transitioning a volume SnapMirror relationship, the following conditions must be met:
 - The secondary cluster should not have an SVM with the same name as that of the primary SVM.
 - The primary cluster should not have an SVM with the same name as that of the secondary SVM.
 - The name of the source 7-Mode system should not conflict with any of the local SVMs or SVMs that are already peered.

You should not upgrade the cluster to a different ONTAP version during transition.



You can upgrade the cluster to a patch release of the same ONTAP version, if required.

Steps

1. From an administration host, verify that the cluster is reachable by using the cluster-management LIF:

```
ssh username@cluster_mgmt_IP
```

2. Enable SSLv3 or FIPS on the cluster:

If you want to enable...	Enter...
SSLv3	system services web modify -ssl3 -enabled true
FIPS 140-2 compliance	system services web modify -ssl-fips -enabled true

When FIPS 140-2 compliance is enabled, SSLv3 is disabled. ONTAP prevents you from enabling SSLv3 when FIPS 140-2 compliance is enabled. If you enable FIPS 140-2 and then subsequently disable it, SSLv3 remains disabled.



The best practice is to enable FIPS because of the security vulnerabilities in SSLv3.

3. Verify that HTTPS is allowed on the cluster management LIF:

- a. View the firewall policy for the cluster management LIF:

```
network interface show -vserver svm_name -lif cluster_mgmt_lif -fields  
firewall-policy
```

```
cluster1::> network interface show -vserver cluster1 -lif
cluster_mgmt -fields firewall-policy
vserver lif      firewall-policy
-----
cluster1  cluster_mgmt mgmt
```

- b. Verify that the firewall policy associated with the cluster management LIF allows HTTPS access:

system services firewall policy show -policy mgmt

```
cluster1::> system services firewall policy show -policy mgmt
Policy          Service      Action IP-List
-----
mgmt
                dns        allow  0.0.0.0/0, ::/0
                http       allow  0.0.0.0/0, ::/0
                https      allow  0.0.0.0/0, ::/0
                ndmp       allow  0.0.0.0/0, ::/0
                ntp        allow  0.0.0.0/0, ::/0
                rsh        deny   0.0.0.0/0, ::/0
                snmp       allow  0.0.0.0/0, ::/0
                ssh        allow  0.0.0.0/0, ::/0
                telnet     deny   0.0.0.0/0, ::/0
9 entries were displayed.
```

System administration

4. Create an intercluster LIF on each node of the cluster for communication between the cluster and 7-Mode system:

- a. **network interface create -vserver *svm_name* -lif *intercluster_lif* -role *intercluster* -home-node *home_node* -home-port *home_port* -address *ip_address* -netmask *netmask***

```
cluster1::> network interface create -vserver cluster1-01 -lif
intercluster_lif -role intercluster -home-node cluster1-01 -home-port
e0c -address 192.0.2.130 -netmask 255.255.255.0
```

- b. Create a static route.

If you are transitioning to...	Run this command...
ONTAP 9.5 or earlier or clustered Data ONTAP 8.3.x	<p>network route create</p> <pre>cluster1::> network route create -vserver vs0 -destination 0.0.0.0/0 -gateway 10.61.208.1</pre>
Clustered Data ONTAP 8.2.x	<p>network routing-groups route create</p> <pre>cluster1::> network routing- groups route create -vserver cluster1-01 -routing-group i192.0.0.0/18 -destination 0.0.0.0/0 - gateway 192.0.2.129</pre>

- c. Verify that you can use the intercluster LIF to ping the 7-Mode system:

network ping -lif *intercluster_lif* -vserver *svm_name* -destination *remote_inetaddress*

```
cluster1::> network ping -lif intercluster_lif -vserver cluster1
-destination system7mode
system7mode is alive
```

For multipathing, you must have two intercluster LIFs on each node.

[Network and LIF management](#)

Preparing 7-Mode aggregates and volumes for transition

Before transition, you must ensure that the 7-Mode aggregates and volumes are eligible for transition and perform some manual steps before transition. For example, some volume types cannot be transitioned and any 32-bit data must be removed from the 7-Mode systems before transition.

Restrictions for transitioning 7-Mode volumes

You must be aware of certain restrictions for transitioning 7-Mode volumes. Some of the restrictions are due to features that are not supported in ONTAP. For some restrictions, you can perform a corrective action that enables you to continue with the transition.

Volume types

The following types of volumes are not supported for transition:

- Traditional volumes

You can use host-based transition methods to transition traditional volumes.

[NetApp Technical Report 4052: Successfully Transitioning to Clustered Data ONTAP \(Data ONTAP 8.2.x and 8.3\)](#)

- FlexCache volumes

Volume states

Transition is blocked if any of the 7-Mode volumes selected for the transition are in one of the following states:

- Offline
- Restricted
- Inconsistent (`wavl inconsistent`)

Volume with qtrees that belong to a different vFiler unit

You cannot transition volumes with qtrees, where the qtrees are owned by a different vFiler unit than that of the volume. Before transition, you must ensure that each volume and all of its qtrees belong to the same vFiler unit by performing one of the following actions:

- Move the qtrees to the vFiler unit that owns the volume.
- Delete the qtrees.

Inode to parent pathname translation setting

The inode to parent pathname translations must be enabled on each volume. You can enable the parent to pathname translations by turning off the `no_i2p` option:

```
vol options vol_name no_i2p off
```

You do not have to wait for the i2p scan to finish, and you can continue with the transition preparation.

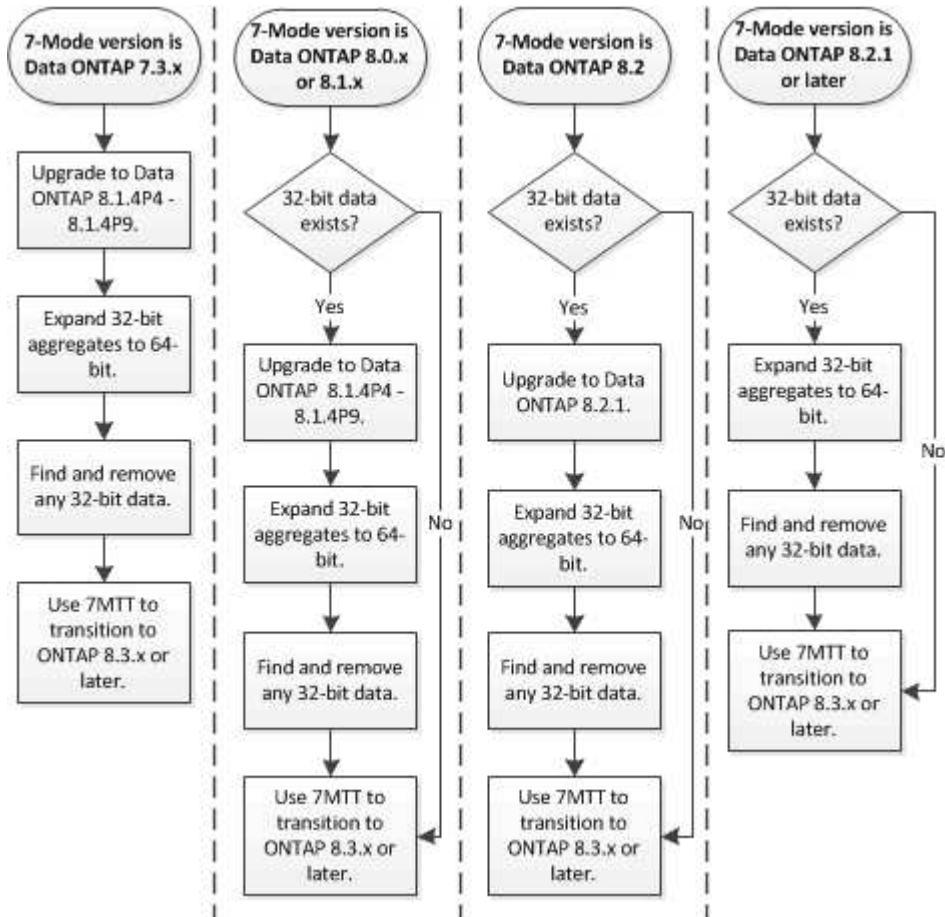
Preparing for transitioning to ONTAP 8.3 and later supported releases

32-bit aggregates, volumes, and Snapshot copies are not supported in ONTAP 8.3 and later. Therefore, you must expand the 32-bit aggregates to 64-bit, and then find and remove any 32-bit volumes and Snapshot copies from the 7-Mode system before transition. Because all 7-Mode versions do not support the capability of expanding 32-bit aggregates and removing 32-bit volumes and Snapshot copies, you might have to upgrade your 7-Mode system before transition.



Clustered Data ONTAP 8.2.x supports 32-bit aggregates, volumes and Snapshot copies. Therefore, you can transition 32-bit data from 7-Mode system to a target cluster running Data ONTAP 8.2.x. However, after the transition, if the target cluster must be upgraded to ONTAP 8.3 or later version, then you must upgrade all the existing 32-bit data on the target cluster to 64-bit format before upgrading the ONTAP version of the target cluster..

You should use the following workflow to decide whether an upgrade is required before transition.



Related information

[NetApp Technical Report 3978: In-Place Expansion of 32-Bit Aggregates to 64-Bit Overview and Best Practices](#)

Expanding an aggregate to the 64-bit format

If your system contains 32-bit aggregates, you must expand them to the 64-bit format on your 7-Mode system *before* transitioning to Data ONTAP 8.3 or later versions, because those versions of Data ONTAP do not support the 32-bit format.

- If the aggregate contains destination volumes for a SnapMirror relationship with a 32-bit source volume, the aggregate containing the source volume must be expanded before expanding the aggregate containing the destination volume.

For volumes in a SnapMirror relationship, the destination volume inherits the format of the source volume while the mirror is intact. If the aggregate you are expanding contains a destination volume whose source is a 32-bit volume and you break the mirror before expanding the aggregate, the destination volume is expanded to the

64-bit format. However, if you reestablish the mirror and the source volume is still 32-bit, the destination volume returns to the 32-bit format. For this reason, you must expand the aggregate containing the source volume before reestablishing the SnapMirror relationship if you want to expand all 32-bit volumes in the aggregate to the 64-bit format.

Steps

- 1. Enter advanced privilege mode:

```
priv set advanced
```

- 2. Initiate the expansion:

```
aggr 64bit-upgrade start aggr_name
```

- 3. Perform the appropriate action:

If the command...	Then...
Initiates successfully	Proceed to the next step.
Indicates that one or more volumes could not be expanded because they did not have enough space	Retry the command, adding the grow-all option.
Indicates that the expansion could not be completed for some other reason	Perform the appropriate action, based on the issue outlined in the error message.

- 4. Display the status of the expansion:

```
aggr 64bit-upgrade status aggr_name
```

The current status of the expansion is displayed. When the message indicates that there is no upgrade in progress, the expansion is complete.

- 5. Confirm that all volumes in the aggregate are 64-bit format:

```
aggr 64bit-upgrade status aggr_name -all
```

- 6. Return to administrative privilege mode:

```
priv set admin
```

The aggregate is expanded to the 64-bit format. However, even if all volumes are expanded, some 32-bit Snapshot copies might remain. The presence of 32-bit Snapshot copies in the source volumes prevents an upgrade or transition to Data ONTAP 8.3 or later.

Finding and removing 32-bit volumes and Snapshot copies

Even if you have expanded all of your aggregates to the 64-bit format, some 32-bit or mixed-format FlexVol volumes or Snapshot copies can remain. These volumes and Snapshot copies must be removed before your data can be accessed by a cluster running Data ONTAP 8.3 or later.

- You must have expanded all 32-bit aggregates on the system to the 64-bit format.

You must repeat the steps in this task for each aggregate that contains 32-bit volumes and Snapshot copies.

Steps

1. Enter advanced mode:

```
priv set advanced
```

2. Display the format of all volumes in the aggregate:

```
aggr 64bit-upgrade status aggr_name -all
```

Each volume in the aggregate is displayed with its format.

3. For each 32-bit or mixed-format volume, determine the reason that the volume has not been expanded to the 64-bit format, and then take the appropriate action.

If you cannot determine the reason that the volume was not expanded, retry the aggregate expansion.

If the volume...	Then...
Is the destination of a SnapMirror relationship	Expand the aggregate containing the source volume to the 64-bit format.
Is a read-only volume (but not a SnapMirror destination)	Make the volume writable and retry the expansion, or destroy the volume.
Did not expand because of insufficient free space in the volume or aggregate	Increase the free space in the volume or aggregate and retry the expansion.

All 32-bit and mixed-format volumes in the aggregate are now 64-bit. You can confirm this by repeating the previous step.

4. Display the format of all Snapshot copies on the system:

```
snap list -fs-block-format
```

5. Remove the 32-bit Snapshot copies by using the `snap delete` command.



This action deletes the data in the Snapshot copies. You must be certain that you do not need to retain the Snapshot copies before you delete them. Alternatively, you can wait for the 32-bit Snapshot copies to be aged out. The amount of time this takes depends on your Snapshot copy schedule.

If a Snapshot copy is the base Snapshot copy for a FlexClone volume, you must split the FlexClone volume from its parent before you can remove the Snapshot copy.

All 32-bit Snapshot copies are removed. You can confirm this by repeating the previous step.

6. Return to the administrative privilege level:


```
priv set admin
```

Considerations for deduplication and compression

When using compression, the source and the destination volumes must belong to a 64-bit aggregate. All compression and deduplication savings on the source volume are retained over the network during transition. After transition, the destination volume inherits all of the compression and deduplication attributes and storage savings from the source volume.

Transitioning deduplicated and compressed data helps in reducing the network bandwidth during transition because of the following reasons:

- Shared blocks are transferred only once.
- Compression is maintained throughout the transfer.
- Compressed and deduplicated data involves smaller transfer sizes as a result of compression and deduplication space savings; therefore, the transfers are completed more quickly.

You should not start compression or deduplication of existing data on the source volume during transition. If deduplication or compression is in progress, you should start the transition only after the deduplication or compression operation is complete. Therefore, undeduplicated or uncompressed data and additional temporary metadata files are not sent over the network to the destination volume.

For deduplication and compression to take effect on any new data written on the ONTAP volume, you must enable deduplication and compression schedules after the transition.

Beginning with Data ONTAP 8.1, deduplication maintains a partially ordered fingerprint database in the volume along with the aggregate copy. As a result, the destination system will have the space savings from the source volume as well as a copy of the ordered fingerprint database. After migration, when volume efficiency is run on the new volume for the first time, the aggregate fingerprint database is automatically constructed from the copy in the destination volume. This can result in a one-time increase in the time it takes for volume efficiency operations to complete.

If your source volume is running a Data ONTAP operating in 7-Mode prior to 8.1, you must run the `volume efficiency start` command with the `-scan-old-data` option parameter to optimize space savings. After the migration is complete, you should verify whether the deduplication schedule meets your requirements on your cluster and consider switching to a volume efficiency policy.

Considerations for FlexClone volumes

When transitioning FlexClone volumes to the SVM, the clones are split from the parent volumes and are transitioned as FlexVol volumes to the destination cluster. As a result, the clone hierarchy and storage efficiency are lost in the transition process.

If the target cluster is running Data ONTAP 8.3 or earlier, FlexClone volumes cannot be created from Snapshot copies that are transitioned from 7-Mode. You can create FlexClone volumes only from new Snapshot copies that are created after the volume is transitioned to ONTAP. Beginning with clustered Data ONTAP 8.3.1, you can create FlexClone volumes from Snapshot copies that are transitioned from 7-Mode.

Considerations for quotas

You must be aware of how quotas are transitioned when “apply configuration” (precutover) is run in read-only and read-write mode.

Quotas are applied in the following ways during the precutover phase:

- Read-only mode

Quotas are not applied in the precutover read-only mode on the ONTAP system; they are applied only during the storage cutover phase.

- Read-write mode

Quotas are applied in the precutover read-write mode on the ONTAP system, so that you can test them in ONTAP. However, the quotas are removed during resynchronization (after testing is completed) of the ONTAP volumes. The quotas are applied again during the storage cutover phase.

Support for transitioning SnapLock volumes

The 7-Mode Transition Tool supports the transition of SnapLock volumes to target clusters running any ONTAP 9.0 release except 9.6.

The SnapLock Enterprise and SnapLock Compliance volumes are supported for transition to target clusters that are running any ONTAP release except 9.6. However, SnapLock Compliance volume transition is not supported to the target clusters that are in MetroCluster configurations.

Considerations for transitioning of SnapLock Enterprise volumes

The 7-Mode Transition Tool supports the transition of stand-alone SnapLock Enterprise volumes and SnapLock Enterprise volumes that are in a SnapMirror relationship.

The workflow for transitioning SnapLock Enterprise volumes is the same as for FlexVol volumes.

SnapMirror relationships are preserved during the transition.



The 7-Mode Transition Tool only supports like-to-like transition for SnapMirror relationships of SnapLock Enterprise volumes. That is, both the source and destination volumes must be SnapLock Enterprise volumes.

Considerations for transitioning of SnapLock Compliance volumes

The 7-Mode Transition Tool supports the transition of standalone SnapLock Compliance volumes and SnapLock Compliance volumes that are in a SnapMirror relationship.

The workflow for transitioning standalone SnapLock Compliance volumes is the same as for transitioning FlexVol volumes.

The transition of SnapMirror relationships for SnapLock Compliance volumes is not automated by the 7-Mode Transition Tool. You must transition the primary and secondary SnapLock Compliance volumes as stand-alone volumes, and then manually resynchronize the relationships.

You can include the SnapLock Compliance volumes (both stand-alone and the volumes that are in SnapMirror relationships) as a standalone volume in stand-alone, primary, and secondary projects.

The precutover read/write mode is not supported for projects with SnapLock Compliance volumes. It is a best practice to create separate projects for SnapLock Compliance volumes and non-SnapLock Compliance volumes because the precutover read/write mode is not supported if SnapLock Compliance volumes are included in the project.

During the cutover operation, if the selected volume is a SnapLock Compliance volume and it is the destination of a SnapMirror relationship, then the SnapMirror relationship between the 7-Mode volume and the ONTAP volume is deleted without SnapMirror break operation. This action enables the secondary ONTAP SnapLock Compliance volumes to remain in read-only mode. The secondary ONTAP SnapLock Compliance volumes must be in read-only mode for the resynchronization operation to be successful between the primary and secondary SnapLock Compliance volumes.

See [How to transition the 7-Mode SnapLock Compliance volumes with SnapMirror relationship to clustered Data ONTAP](#)

Considerations for transitioning of SnapLock Audit volumes

The 7-Mode Transition Tool supports the transition of SnapLock Audit volumes. The workflow to transition SnapLock Audit volumes is the same as the transition of SnapLock Compliance volumes.

After you transition audit volumes to the ONTAP, you must manually designate the transitioned audit volume as SnapLock Audit volume for the target SVM.

In ONTAP, the audit volumes are configured at an SVM level. In Data ONTAP operating in 7-Mode, an audit volume serves as a consolidated repository for all of the volumes in the controller across the vFiler units.

SnapLock Audit volumes are a type of SnapLock Compliance volume. The transition of SnapLock Audit volumes is not supported if the target cluster is in a MetroCluster configuration.

See [How to configure audit volume in clustered Data ONTAP for the transitioned SnapLock volumes](#)

Considerations for transitioning of 7-Mode SnapLock options

The 7-Mode Transition Tool supports the transition of a few 7-Mode options that are related to SnapLock volumes.

Data ONTAP operating in 7-Mode has the following options that are related to SnapLock volumes:

- `snaplock.autocommit_period`

This option is at a volume level in ONTAP, and is transitioned to ONTAP during the transition.

- `snaplock.compliance.write_verify`

This option is not applicable in ONTAP.

- `snaplock.log.default_retention`
- `snaplock.log.maximum_size`

Although the `snaplock.log.default_retention` and `snaplock.log.maximum_size` options are supported in ONTAP, the settings configured in these options are not transitioned by the 7-Mode Transition Tool. You must manually set these options for audit volumes after the transition is completed.

Considerations for using Chain of Custody verification for 7-Mode SnapLock volumes

You should be aware of the considerations for using Chain of Custody verification for 7-Mode SnapLock volumes.

- The SnapLock Chain of Custody verification must be performed only if it is a requirement for the transition of SnapLock volumes.

You can perform the Chain of Custody verification for all or a subset of SnapLock volumes in the project.

- The SnapLock Chain of Custody verification can take a significant amount of time based on the number of files on the 7-Mode SnapLock volumes.
- The Chain of Custody verification is supported only for read/write 7-Mode SnapLock volumes

The Chain of Custody verification is not supported for read-only volumes.

- The Chain of Custody verification is not supported for SnapLock volumes containing files that have names with non-ASCII characters.

Preparing to transition name services

Name service configurations that include DNS, LDAP, NIS, hosts, name services switch, UNIX users and groups, and netgroups configurations are transitioned by the 7-Mode Transition Tool. You must be aware of some considerations before transitioning name services configurations.

Name services transition: supported and unsupported configurations, and required manual steps

You must be aware of the name services configurations that are transitioned by the 7-Mode Transition Tool. Some name services configurations are not transitioned to ONTAP because either these are not supported in ONTAP or these must be manually transitioned.

You should verify all the precheck error and warning messages to evaluate the impact of such configurations on transition.

Configurations that are transitioned

At a high level, the following name services configurations are transitioned by the 7-Mode Transition Tool:

- DNS configuration (`/etc/resolv.conf`)
- LDAP configuration
- NIS configuration

- Name service switch configuration (`/etc/nsswitch.conf` and `/etc/resolv.conf`)
- Hosts configuration (`/etc/hosts`)
- UNIX users and groups (`/etc/passwd` and `/etc/group`)
- Netgroups configuration (`/etc/netgroup`)

See the precheck results for details about these name services configurations.

Unsupported configurations in ONTAP

- NIS slave
- NIS broadcast
- NIS groups caching
- Dynamic DNS
- DNS cache
- Shadow database
- Host database sources other than file or DNS

ONTAP supports only file and DNS for host lookup; other database sources are not supported. Host lookup order in the `/etc/nsswitch.conf` is ignored during transition.

Configurations that must be manually configured

You must manually configure the following LDAP options on the SVMs:

- `ldap.usermap.attribute.unixaccount`
- `ldap.password`
- `ldap.usermap.base`
- `ldap.ssl.enable`

Related information

[NFS management](#)

[Network and LIF management](#)

Considerations for transitioning DNS, NIS, and LDAP configurations

You should be aware of how the DNS, NIS, and LDAP configurations in Data ONTAP operating in 7-Mode are transitioned and applied in ONTAP.

Considerations for DNS transition

For DNS configurations, a maximum of six domain names and three name servers per SVM are supported in ONTAP. If the unique number of domain names or name servers across 7-Mode systems and the target SVM exceed the supported limit, the 7-Mode Transition Tool reports a blocking error. To continue with the transition, you should ignore the transition of the DNS configuration from the tool.



If you ignore the transition of the DNS configuration, you must manually configure DNS on the target SVM.

Considerations for NIS transition

- The length of the NIS domain name on the 7-Mode system must not exceed 64 characters.
- For transitioning to target cluster versions running ONTAP 9.1 or earlier, the `nis.servers` option on the 7-Mode system must be configured only with IP addresses, and not a fully qualified domain name (FQDN).

You must configure the `nis.servers` option on the 7-Mode system with IP addresses before transition if you are transitioning to a cluster running ONTAP 9.1 or earlier. Transition is supported if you have the `nis.servers` option on the 7-Mode system configured with an FQDN and you are transitioning to a cluster running any version of ONTAP between 9.2 and 9.5.

Considerations for LDAP transition

- If multiple base values and scope values are set for the `ldap.base`, `ldap.base.passwd`, `ldap.base.group`, or `ldap.base.netgroup` option, and if you are transitioning to clustered Data ONTAP 8.2 or 8.2.1, only one value for each option is transitioned.

After transition, there might be lookup issues for these options. You must manually add the base values and scope values after transition.

- If multiple scope values are set for the `ldap.base`, `ldap.base.passwd`, `ldap.base.group`, or `ldap.base.netgroup` option, and if you are transitioning to clustered Data ONTAP 8.2.2, only one value for each option is transitioned.
- If separate base values and scope values are specified for user mapping (`ldap.usermap.base`) and user password (`ldap.base.passwd`) lookups in the 7-Mode system, the base values and scope values for only the user password are transitioned.

The base values and scope values are used for user mapping and user password lookups in ONTAP, which can cause security issues. You must manually add the base values and scope values for user mapping to the user distinguished name (DN) option in ONTAP after transition, if required.

Considerations for transitioning netgroups and UNIX users and groups

Netgroup configuration is transitioned only if the 7-Mode `/etc/netgroup` file is less than 5 MB in size. UNIX users and groups are transitioned only if the total number of UNIX users and groups on the SVM do not exceed the limits for users and groups in ONTAP.

Considerations for netgroups

If the `/etc/netgroup` file on 7-Mode is greater than 5 MB, the netgroup configuration is not transitioned. You must perform one of the following actions to continue with the transition:

- Exclude the transition of netgroups.
- Move the netgroup configuration to NIS or LDAP servers before transition.

Considerations for UNIX users and groups

If the total number of transitioning UNIX users and groups exceed the limit of UNIX users and groups in ONTAP, the 7-Mode Transition Tool blocks the transition. You must perform one of the following actions to continue with the transition:

- Exclude the transition of UNIX users and groups.
- Move the UNIX users and groups to NIS or LDAP servers before transition.

Related information

[NFS management](#)

Preparing for NFS transition

If NFS is licensed and NFS service is running on the systems operating in 7-Mode, you must manually prepare the cluster and target SVM for transitioning NFS configurations. You must also be aware of what configurations are transitioned.

Some NFS configurations operating in 7-Mode are not supported in ONTAP. Some configurations are not transitioned by the 7-Mode Transition Tool and must be manually applied to the SVM.

Prerequisites for transitioning NFS configurations

NFS configurations are transitioned by the 7-Mode Transition Tool only when certain prerequisites are met on the 7-Mode system and the cluster. If any of the conditions are not met, the tool does not transition the configuration.

7-Mode prerequisites

- NFS must be licensed.
- If MultiStore is licensed, NFS must be enabled on the vFiler unit that owns the transitioning volumes.
- For transitioning a Microsoft Active Directory (AD) based Kerberos server to a new SVM, a DNS entry must exist for the AD domain.



To transition the Kerberos configuration, at least one LIF must be transitioned as part of the project and the LIF must be resolvable to a host name.

- If you want to transition in-memory export rules, you must add them to the `/etc/exports` file before transition.

The 7-Mode Transition Tool transitions only the persistent export rules that are defined in the `/etc/exports` file.

Cluster prerequisites

- NFS must be licensed.
- For transitioning a Microsoft AD-based Kerberos server to an existing SVM with DNS configured, a DNS entry must exist for the AD domain.

- The clock skew between the Kerberos key distribution center (KDC) and the ONTAP system must be less than or equal to 5 minutes.

Related information

[How NFS exports are transitioned](#)

[NetApp Documentation: ONTAP 9](#)

NFS transition: supported and unsupported configurations, and required manual steps

Some NFS configurations are not transitioned to ONTAP because they are not supported in ONTAP, there are functionality differences from 7-Mode, or they must be manually transitioned. You should verify all of the precheck errors and warning messages to evaluate the impact of such configurations on transition.

Supported configurations for transition

At a high level, the following NFS configurations are transitioned by the 7-Mode Transition Tool:

- NFS options:
 - `nfs.udp.xfersize`
 - `nfs.v4.id.domain`
 - `nfs.v4.acl.max.aces`
 - `nfs.tcp.xfersize`
 - `nfs.rpcsec.ctx.high`
 - `nfs.rpcsec.ctx.idle`
 - `nfs.response.trigger`
 - `wapl.default_nt_user`
 - `nfs.mount_rotonly`
 - `nfs.tcp.enable`
 - `nfs.udp.enable`
 - `nfs.response.trace`
 - `nfs.v4.read_delegation`
 - `nfs.v4.write_delegation`
 - `nfs.v4.acl.enable`
 - `nfs.vstorage.enable`
 - `nfs.v3.enable`
 - `nfs.v4.enable`
- NFS export rule:

If the export rule is configured with the `-actual` option, the exported path (alias path) is ignored and the export rule is configured with the actual path.

- Export rules with Kerberos security `krb5p`
- Kerberos configuration

See the precheck results for details about these NFS configurations.

Unsupported configurations in ONTAP

The following NFS configurations are not supported in ONTAP:

- Subvolume NFS exports other than `qtree`-level NFS exports
- WebNFS
- PC-NFS
- NFSv2
- Fencing of NFS clients from one or more file system paths
- Some NFS options

See the precheck warning messages for a complete list of unsupported options.

Configurations that must be manually transitioned

There are some NFS configurations that are supported in ONTAP, but are not transitioned by the 7-Mode Transition Tool.

The following NFS configurations generate a warning message in the precheck operation, and you must manually apply the configurations on the SVM:

- NFS audit configuration
- NFS options:
 - `rpc.nsm.tcp.port`
 - `rpc.nsm.udp.port`
 - `rpc.mountd.tcp.port`
 - `rpc.mountd.udp.port`
 - `nfs.export.neg.timeout`
 - `nfs.export.pos.timeout`
 - `nfs.export.harvest.timeout` Use the `vserver nfs modify` command to modify the configuration of an NFS-enabled storage virtual machine (SVM).
- Export rules with Kerberos security `krb5p`

Configurations that are functionally different in ONTAP

The following NFS configurations are functionally different in ONTAP:

- NFS export rules

- NFS export access cache
- NFS diagnostic commands
- Support for the `showmount` command
- NFS Kerberos encryption
- NLM version support

Related information

[NFS management](#)

How NFS exports are transitioned

You must be aware of how NFS exports are configured on the SVM after transition. You might have to perform some manual steps if the 7-Mode export configurations are not supported in ONTAP.

You must be aware of the following considerations about NFS exports transition:

- If the SVM root volume is not exported to allow read-only access to all NFS clients, the 7-Mode Transition Tool creates a new export policy that allows read-only access for all the NFS clients and exports the root volume of the SVM with the new export policy.

To ensure that all the transitioned volumes or qtrees are mountable, the root volume of the SVM must be allowed read-only access for all the NFS clients.

- When 7-Mode volumes with export configurations that are not supported in ONTAP are transitioned, these volumes are exported to allow read-only permissions to all NFS clients on the SVM.

Export policies for these volumes must be configured manually after transition to provide the required access permissions.

- When 7-Mode qtrees with export configurations that are not supported in ONTAP are transitioned, they inherit the export policy of the parent volume.

Export policies for these qtrees must be configured manually after transition to provide the required access permissions.

- In ONTAP, for an NFS client to mount a qtree, the NFS client must have read-only permissions at all the parent junction paths up to the SVM's root volume junction path (that is, /).

For NFS clients to mount qtrees, the qtrees must belong to a volume that has read-only permission. Without the read-only permissions at the volume level, the NFS clients cannot mount the qtree.

- If the same host is specified in the combination of read-only, read-write, and root access permission lists, you must evaluate the transitioned export rules after transition to determine appropriate access privilege for the hosts.

[NetApp Technical Report 4067: NFS Best Practice and Implementation Guide](#)

Example: Modifying the export policy of a volume to allow access to a qtree

Consider the following export rule configured in the 7-Mode storage system (192.168.26.18) that allows

read/write access to the volume volstd10 and qtree qtree1 for the NFS client 192.168.10.10:

```
/vol/volstd10/qtree1 -sec=sys,rw=192.168.10.10,nosuid
/vol/volstd10 -sec=sys,rw=192.168.11.11,nosuid
```

After transition, the export policy of the volume volstd10 in ONTAP is as shown below:

```
cluster-01::> export-policy rule show -vserver std_22 -policyname std_2226
-instance
(vserver export-policy rule show)

Vserver: std_22
Policy Name: std_2226
Rule Index: 1
Access Protocol: any
Client Match Hostname, IP Address, Netgroup, or Domain: 192.168.11.11
RO Access Rule: sys
RW Access Rule: sys
User ID To Which Anonymous Users Are Mapped: 65534
Superuser Security Types: none
Honor SetUID Bits in SETATTR: false
Allow Creation of Devices: true

cluster-01::>
```

After transition, the export policy of the qtree qtree1 in ONTAP is as shown below:

```
cluster-01::> export-policy rule show -vserver std_22 -policyname
std_2225 -instance
(vserver export-policy rule show)

Vserver: std_22
Policy Name: std_2225
Rule Index: 1
Access Protocol: any
Client Match Hostname, IP Address, Netgroup, or Domain: 192.168.10.10
RO Access Rule: sys
RW Access Rule: sys
User ID To Which Anonymous Users Are Mapped: 65534
Superuser Security Types: none
Honor SetUID Bits in SETATTR: false
Allow Creation of Devices: true

cluster-01::>
```

For the NFS client 192.168.10.10 to access the qtree, the NFS client 192.168.10.10 must have read-only

access to the qtree's parent volume.

The following output shows that the NFS client is denied access while mounting the qtree:

```
[root@192.168.10.10 ]# mount 192.168.35.223:/vol/volstd10/qtree1
transition_volume_qtreemount:192.168.35.223:/vol/volstd10/qtree1 failed,
reason
given by server: Permission denied [root@192.168.10.10 ]#
```

You must manually modify the export policy of the volume to provide read-only access to the NFS client 192.168.10.10.

```

cluster-01::> export-policy rule create -vserver std_22 -policyname
std_2226 -clientmatch
192.168.10.10 -rorule sys -rwrule never -allow-suid false -allow-dev true
-superuser none -protocol nfs
    (vserver export-policy rule create)

cluster-01::> export-policy rule show -vserver std_22 -policyname std_2226
-instance
    (vserver export-policy rule show)

                                Vserver: std_22
                                Policy Name: std_2226
                                Rule Index: 1
                                Access Protocol: any
Client Match Hostname, IP Address, Netgroup, or Domain: 192.168.11.11
                                RO Access Rule: sys
                                RW Access Rule: sys
User ID To Which Anonymous Users Are Mapped: 65534
                                Superuser Security Types: none
                                Honor SetUID Bits in SETATTR: false
                                Allow Creation of Devices: true

**                                Vserver: std_22
                                Policy Name: std_2226
                                Rule Index: 2
                                Access Protocol: nfs
Client Match Hostname, IP Address, Netgroup, or Domain: 192.168.10.10
                                RO Access Rule: sys
                                RW Access Rule: never
User ID To Which Anonymous Users Are Mapped: 65534
                                Superuser Security Types: none
                                Honor SetUID Bits in SETATTR: false
                                Allow Creation of Devices: true**

cluster-01::>

```

Example: How qtree export rules differ in 7-Mode and ONTAP

In the 7-Mode storage system, when an NFS client accesses a qtree through the mount point of its parent volume, the qtree export rules are ignored and the export rules of its parent volume are in effect. However, in ONTAP, qtree export rules are always enforced whether NFS client mounts to the qtree directly or it accesses the qtree through the mount point of its parent volume. This example is specifically applicable for NFSv4.

The following is an example of an export rule on the 7-Mode storage system (192.168.26.18):

```
/vol/volstd10/qtrees1 -sec=sys,ro=192.168.10.10,nosuid  
/vol/volstd10 -sec=sys,rw=192.168.10.10,nosuid
```

On the 7-Mode storage system, the NFS client 192.168.10.10 has only read-only access to the qtrees. However, when the client accesses the qtrees through the mount point of its parent volume, the client can write to the qtrees because the client has read/write access to the volume.

```
[root@192.168.10.10]# mount 192.168.26.18:/vol/volstd10 transition_volume  
[root@192.168.10.10]# cd transition_volume/qtrees1  
[root@192.168.10.10]# ls transition_volume/qtrees1  
[root@192.168.10.10]# mkdir new_folder  
[root@192.168.10.10]# ls  
new_folder  
[root@192.168.10.10]#
```

In ONTAP, the NFS client 192.168.10.10 has only read-only access to the qtrees qtrees1 when the client accesses the qtrees directly or through the mount point of the qtrees's parent volume.

After transition, you must evaluate the impact of enforcing the NFS export policies, and if necessary modify the processes to the new way of enforcing NFS export policies in ONTAP.

Related information

[NFS management](#)

Preparing for SMB/CIFS transition

If SMB/CIFS is licensed and SMB/CIFS service is running on the 7-Mode systems, you must manually perform some tasks, such as adding the SMB/CIFS license and creating a SMB/CIFS server, on the target cluster and SVM for transitioning SMB/CIFS configurations.

You must also be aware of what configurations are transitioned. Some SMB/CIFS configurations operating in 7-Mode are not supported in ONTAP. Some configurations are not transitioned by the 7-Mode Transition Tool and must be manually applied to the SVM.

Prerequisites for transitioning CIFS configurations

CIFS configurations are transitioned by the 7-Mode Transition Tool only when certain prerequisites are met on the 7-Mode system and cluster. If any of the conditions are not met, the tool does not transition the configuration.

7-Mode prerequisites

- The CIFS license must be added.
- If the MultiStore license is enabled, CIFS must be added to the list of allowed protocols for the vFiler unit that owns the transitioning volumes.

- CIFS must be set up and running during transition.
- The authentication type for CIFS must be Active Directory (AD) or Workgroup.

Cluster prerequisites

- The CIFS license must be added.
- CIFS must be added to the list of allowed protocols for the SVM.
- DNS must be configured for the SVM.
- The following CIFS authentication methods are supported in different ONTAP versions:
 - Clustered Data ONTAP 8.2.x and 8.3.x support AD authentication.
 - ONTAP 9.0 or later supports AD authentication and Workgroup authentication.
- Use the following table to decide which authentication must be used on the target SVM:

7-Mode authentication method	Clustered Data ONTAP 8.2.x and 8.3.x authentication method	ONTAP 9.5 or earlier authentication method
AD	AD	AD
Workgroup	AD	Workgroup or AD

- You can transition CIFS configuration from 7-Mode to ONTAP if the AD domains does not match between 7-Mode CIFS server and target SVM CIFS server. The tool triggers an ignorable blocking error when an AD domain name mismatch is detected. To proceed with the transition, acknowledge the blocking error.
- The CIFS server must be manually configured before the apply configuration phase (precutover).

You can create the CIFS server on the SVM in the following two ways:

If you want to...	Do the following...
<p>Transfer or preserve the CIFS server identity to the target SVM</p>	<ul style="list-style-type: none"> You must plan to transition all of the volumes in the source 7-Mode system or vFiler unit in a single project. <p>This plan is required because the 7-Mode system loses the original CIFS server identity after the transition and cannot serve clients. The maximum number of volumes that can be transitioned in one project is 160; therefore, to preserve the CIFS server identity, the 7-Mode system can have a maximum of 160 volumes and all these volumes must be transitioned in a single project.</p> <p>You have the following two options to create the CIFS server:</p> <ol style="list-style-type: none"> Applicable for all versions of ONTAP: <ul style="list-style-type: none"> Before the “apply configuration” phase (precutover), you must reconfigure the CIFS server on the 7-Mode system by using a temporary CIFS identity. <p>This reconfiguration allows the original CIFS server identity to be configured on the SVM. You must verify that the CIFS server is running on the 7-Mode system during the “apply configuration” phase (precutover) operation with the new temporary identity. This action is required to read CIFS configurations from 7-Mode during precutover.</p> <ul style="list-style-type: none"> You must configure the CIFS server on the target SVM with the original 7-Mode CIFS identity. After these conditions are met, you can perform the precutover operation. <p>You must then plan to perform the storage cutover immediately after precutover for enabling client access to ONTAP volumes.</p> Applicable for ONTAP releases 9.0 through 9.5: <ul style="list-style-type: none"> Use the <code>vserver cifs modify</code> command to change the CIFS server name (CIFS Server NetBIOS Name). <p>Using this feature, you should create a CIFS server on the target SVM with a temporary identity, and then perform the apply configuration (precutover) operation.</p>

If you want to...	Do the following...
Use a new identity	<ul style="list-style-type: none"> • Before the “apply configuration” phase (precutover), you must configure the CIFS server on the target SVM with a new CIFS identity. • You must verify that the CIFS server is up and running on the 7-Mode system during the “apply configuration” phase (precutover) operation. <p>This action is required to read CIFS configurations from 7-Mode systems during the “apply configuration” phase (precutover).</p> <p>After these conditions are met, you can perform the precutover operation. You can then test the SVM configurations and plan to perform the storage cutover operation.</p>

Related information

[Considerations for transitioning CIFS local users and groups](#)

Supported and unsupported CIFS configurations for transition to ONTAP

Some CIFS configurations are not transitioned to ONTAP because either they are not supported in ONTAP or they must be manually transitioned. You should verify all precheck error and warning messages to evaluate the impact of such configurations on transition.

Configurations that are supported for transition

At a high level, the 7-Mode Transition Tool transitions the following CIFS configurations:

- CIFS preferred DC configuration
- User mapping configuration:
 - `/etc/usermap.cfg`
 - `waf1.nt_admin_priv_map_to_root`
- CIFS local users and groups
- Symlink and widelink configuration (`/etc/symlink.translations`)
- CIFS audit configuration
- CIFS shares
- CIFS share ACLs
- CIFS home directory configuration
- CIFS options:

- `cifs.gpo.enable`
- `cifs.smb2.enable`
- `cifs.smb2.signing.required`
- `cifs.wins_servers`
- `cifs.grant_implicit_exe_perms`
- `cifs.restrict_anonymous`
- SMB2 connections to external servers, such as a domain controller. The following command implements this support:
 - **`cifs security modify -vserver SVM1 -smb2-enabled-for-dc-connections`**
- FPolicy native file blocking configuration

See the precheck results for details about these CIFS configurations.

Configurations that are not supported in ONTAP

The following 7-Mode configurations are not supported in ONTAP. Therefore, these configurations cannot be transitioned.

- NT4, and password authentication types
- Separate options for SMB1 and SMB2 signing
- CIFS statistics on a per-client basis
- ◦ Authentication for clients earlier than Windows NT
- Auditing of account management events for local users and groups
- Usermap entries with IP addresses, host names, network names, or network names with subnet specified in dotted notation
- CIFS shares with access restriction for machine accounts

Machine accounts can access all shares after transition.

Configurations that must be manually transitioned

Some CIFS configurations are supported in ONTAP, but are not transitioned by the 7-Mode Transition Tool.

The following CIFS configurations generate a warning message in the precheck. You must manually apply these configurations on the SVM:

- Antivirus settings
- FPolicy configurations

7-Mode FPolicy and antivirus servers do not work with ONTAP. You must contact the server vendors for upgrading these servers. However, you must not decommission the 7-Mode FPolicy and antivirus servers until you commit the transition. These are required in case you decide to roll back the transition.

- BranchCache configurations
- Character mapping configuration (charmap)

- Forcegroup attribute of CIFS shares to create files with a specified UNIX group as owning group
- Maxusers attribute of CIFS shares to specify the maximum number of simultaneous connections allowed to a 7-Mode CIFS share
- Storage-Level Access Guard (SLAG) configurations
- Share-level ACLs with UNIX-style permission
- Share ACLs for UNIX users and groups
- LAN Manager authentication level
- NetBIOS aliases
- CIFS search domains
- Some CIFS options

See the precheck results for details about these options.

Considerations for transitioning CIFS local users and groups

You must be aware of the considerations for running the transition operations when migrating CIFS local users and groups.

- If the destination cluster is running clustered Data ONTAP 8.2, transition should not be attempted for 7-Mode volumes that are serving CIFS data and are being accessed by local users and groups.

The 7-Mode Transition Tool does not support the transition of local users and groups to clustered Data ONTAP 8.2.

- Transition of CIFS data-serving volumes from a 7-Mode controller or a vFiler unit that has local users and groups to an SVM that has non-BUILTIN CIFS local users and groups is not supported.

The SVM must have only BUILTIN CIFS local users and groups for transition.

While transitioning local users and groups from a specific 7-Mode controller or a vFiler unit to a specific SVM, local users and groups from the first transition project are transitioned. In the subsequent transition of projects with the same 7-Mode controller or vFiler unit to the same SVM, the transition of local users and groups is ignored, although transition succeeds. The local user's name on the 7-Mode system must not be the same as the CIFS server name on the SVM.

- You must be aware of the limits on the number of local users and groups supported in clustered Data ONTAP 8.2.1 and later.
- A local user account with an empty password or local user accounts with passwords containing more than 14 characters on the 7-Mode system are transitioned to ONTAP software with the password **cifsUser@1**.

After the transition is complete, you can access these users from the Windows system by using the password **cifsUser@1**. You must then manually change the password for such CIFS local users on the SVM by using the following command:

```
cifs users-and-groups local-user set-password -vserver svm_name -user-name user_name.
```

- If the 7-Mode Transition Tool IP address is not reachable from the target ONTAP software, the 7-Mode Transition Tool blocks the transition of CIFS local users and groups to the ONTAP software during the

precheck phase. If you see this error during the precheck phase, use the

network ping -node local -destination *ip_address*

command to make sure the 7-Mode Transition Tool IP address is reachable from the target ONTAP software. You can edit the `\etc\conf\transition-tool.conf` file that is installed with the 7-Mode Transition Tool to modify any configuration option that is used by the tool, such as the 7-Mode Transition Tool IP address.

- The SVM to which the local users and groups are transitioned must have a data LIF.
- If a local group has multiple member system identifiers (SIDs) mapped to a single domain user or group on the 7-Mode system, the 7-Mode Transition Tool blocks the transition of local users and groups to ONTAP during the precheck phase.

If you see this error during the precheck phase, you must manually remove the additional SIDs that are mapped to a single domain user or group on the 7-Mode system. You must then rerun the precheck operation with only a single SID mapped to the domain user or group.

[Troubleshooting Workflow: CIFS: Device attached to the system is not functioning](#)

Related information

[SMB/CIFS management](#)

Preparing for MetroCluster configuration transition

Before transitioning to a MetroCluster configuration, you must understand the requirements and considerations for transitioning 7-Mode volumes to a MetroCluster configuration in ONTAP.

Prerequisites

- The MetroCluster configuration in ONTAP must already be set up.
- The SVM type must be `sync-source`.
- The 7-Mode controllers must not be in a taken over state or waiting for a giveback.
- The nodes in the MetroCluster configuration in ONTAP must not be switched over or waiting for a switchback.

Considerations

- Transitioning SnapLock Compliance volumes is not supported if the target cluster is in a MetroCluster configuration.
- You can transition volumes from a 7-Mode controller, HA configuration, or MetroCluster configuration to a MetroCluster configuration in ONTAP as stand-alone volumes.
- If a 7-Mode MetroCluster configuration has volumes that are in volume SnapMirror relationships with volumes in another 7-Mode controller, you can transition the SnapMirror relationships as primary and secondary relationships.

You should install the 7-Mode Transition Tool on each MetroCluster site and transition the volumes from each site.

- Different subnets configured for a 7-Mode fabric MetroCluster configuration cannot be configured on the MetroCluster configuration in ONTAP.
- The preferred port configured in a 7-Mode fabric MetroCluster configuration cannot be configured for the MetroCluster configurations in ONTAP.
- If your 7-Mode fabric MetroCluster configuration is using Brocade 6510 switches, you can share the existing switch fabrics with the new MetroCluster configuration in ONTAP.

It is best to share the switch fabrics only for the duration of the transition.

[Fabric-attached MetroCluster installation and configuration, ONTAP 9.8 or earlier](#)

[Fabric-attached MetroCluster installation and configuration, ONTAP 9.9.1](#)

- The cron job schedules created during transition are not replicated to the remote site, and therefore the negotiated switchover fails after transition.

You must manually create the cron job schedules on the remote site after the transition.

Related information

[Configuring cron job schedules on the remote site after transitioning a MetroCluster configuration](#)

[Impact of takeover and giveback on transition](#)

[Transitioning a MetroCluster configuration that failed due to switchover or switchback](#)

Preparing for SAN transition

Before transitioning a SAN environment, you must understand what configurations are supported for SAN transition, create SAN LIFs on the SVM, and prepare the SAN hosts for transition.

Preparing SAN hosts for transition

Before transitioning a SAN environment, you must perform some manual steps to prepare the SAN hosts for transition.

You must have generated the inventory workbook for the SAN hosts by using the Inventory Collect Tool.

[Host and storage transition information collection](#)

Steps

1. Verify that the host is supported for transition.

[NetApp Interoperability Matrix Tool](#)

2. Perform the pretransition steps on the host.

[SAN host transition and remediation](#)

Configuring zones by using the FC zone plan

Before transitioning a SAN FC environment, you must configure zones by using the FC zone planner to group the initiator hosts and targets.

- The cluster and initiator hosts must be connected to the switch.
- The FC zone script file must be accessible.

Steps

1. If there are any changes to the igroup configurations on the 7-Mode systems, modify and regenerate the FC zone plan.

[Generating an assessment report by adding systems to the 7-Mode Transition Tool](#)

2. Log in to the CLI of the switch.
3. Copy and execute the required zone commands one at a time.

The following example runs the zone commands on the switch:

```
switch1:admin>config terminal
# Enable NPIV feature
feature npiv
zone name auto_transition_igroup_d31_194bf3 vsan 10
member pwn 21:00:00:c0:dd:19:4b:f3
member pwn 20:07:00:a0:98:32:99:07
member pwn 20:09:00:a0:98:32:99:07
.....
.....
.....
copy running-config startup-config
```

4. Verify the data access from the cluster by using the test initiator hosts.
5. After the verification is complete, perform the following steps:
 - a. Disconnect the test initiator hosts.
 - b. Remove the zone configuration.

Creating SAN LIFs before transition

Because FC and iSCSI LIFs are not transitioned by the 7-Mode Transition Tool, you must create these LIFs on the SVMs before transition. You must configure SAN LIFs on both the nodes that own the LUN and the node's HA partner.

The required SAN (FC or iSCSI) license must be added to the cluster.

For redundancy, you must create SAN LIFs on both the node hosting the LUNs and its HA partner.

Steps

1. Create an FC or iSCSI LIF on the target node to which the LUNs are transitioned, depending on the protocol used:

network interface create

If you want to reuse the 7-Mode IP address for iSCSI LIFs, you must create the LIFs in administrative down state. You can bring these LIFs to the administrative up state after the cutover operation.

2. Create a LIF on the HA partner of the node.
3. Verify that you have set up your LIFs correctly:

network interface show

Related information

[SAN administration](#)


SAN transition: supported and unsupported configurations, and required manual steps

You must be aware of the SAN configurations that are transitioned by the 7-Mode Transition Tool. You should also be aware of the 7-Mode SAN features that are not supported in ONTAP, so that you can take any necessary actions before the transition.

You should verify all of the precheck error and warning messages to evaluate the impact of such configurations on transition.

Configurations that are transitioned

The following SAN configurations are transitioned by the 7-Mode Transition Tool:

- FC and iSCSI services
 - igroups and LUN maps
- 
 - 7-Mode igroups that are not mapped to any LUNs are not transitioned to the target SVMs.
 - For clustered Data ONTAP 8.3.0 and 8.3.1, the transition of igroups and LUN mapping configurations is not supported during the precutover operation.

Instead, the required igroups are created during the cutover operation. For primary and stand-alone volumes, LUNs are mapped to igroups during the cutover operation. However, for secondary volumes, the mapping of LUNs to igroups is not supported during the cutover operation. You must manually map the secondary LUNs after completing the transition of primary volumes.
 - For ONTAP 8.3.2 and later supported releases, igroups and LUN mapping configurations are applied during the precutover operation.

Unsupported configurations in ONTAP

The unsupported configurations in ONTAP are as follows:

- 7-Mode Snapshot copy-backed LUN clones

Snapshot copy-backed LUN clones present in the Snapshot copies are not supported for any restore operation. These LUNs are not accessible in ONTAP. You must split or delete the 7-Mode Snapshot copy-backed LUN clones before transition.

- LUNs with an `ostype` parameter value of `vld`, `image`, or any user-defined string

You must either change the value of the `ostype` parameter for such LUNs or delete the LUNs before transition.

- LUN clone split

You must either wait for the active LUN clone split operations to finish or abort the LUN clone split and delete the LUN before transition.

The following 7-Mode features enable you to continue with the transition process, but are not supported in ONTAP:

- The `lun share` command

Sharing a LUN over NAS protocols

- SnapValidator

Configurations that must be manually transitioned

The following configurations must be transitioned manually:

- SAN LIFs

You must manually create the LIFs before transition.

- Portsets

You must manually configure igroups that are bound to a portset after transition.

- iSCSI access list information
- iSNS configuration
- iSCSI CHAP and RADIUS configurations

Related information

[NFS management](#)

[Network and LIF management](#)

Space considerations when transitioning SAN volumes

You must ensure that sufficient space is available in the volumes during transition. In addition to the space required for storing data and Snapshot copies, the transition process also requires 1 MB of space per LUN for updating certain filesystem metadata.

Before cutover, you can use the `df -h` command on the 7-Mode volume to verify whether free space of 1 MB per LUN is available in the volume. The volume should also have free space equivalent to the amount of data that is expected to be written to the volume before final cutover. If the volume does not have sufficient free space available, the required amount of space must be added to the 7-Mode volume.

If the transition of LUNs fails due to lack of space on the destination volume, the following EMS message is generated: `LUN.vol.proc.fail.no.space: Processing for LUNs in volume vol1 failed due to lack of space.`

In this case, you must set the `filesystem-size-fixed` attribute to `false` on the destination volume, and then add 1 MB per LUN of free space to the volume.

If there are volumes containing space-reserved LUNs, growing the volume by 1MB per LUN might not provide sufficient space. In such cases, the amount of additional space that has to be added is the size of the Snapshot reserve for the volume. After space is added to the destination volume, you can use the `lun transition start` command to transition the LUNs.

Related information

[NetApp Documentation: ONTAP 9](#)

Preparing data protection features for transition

You must perform some manual steps for transitioning 7-Mode SnapMirror relationships. You must also be aware of the data protection relationships that are supported and unsupported for transition.

Data protection transition: supported and unsupported configurations

You can transition a volume that is part of a SnapMirror relationship. However, some data protection and disaster recovery configurations are not supported for transition and therefore you have to perform some manual steps for transitioning these configurations.

Supported configurations

You can transition volume SnapMirror relationships by using the 7-Mode Transition Tool. You can also transition 7-Mode volumes from a MetroCluster Configuration to a MetroCluster Configuration in ONTAP 8.3 and later supported releases.

Unsupported configurations

- SnapVault relationships

Volumes that are the source of a SnapVault relationship can be migrated; however, the SnapVault relationship is not transitioned. A volume that is the destination of a SnapVault relationship can be migrated

only after the SnapVault backups are stopped.

[NetApp Technical Report 4052: Successfully Transitioning to Clustered Data ONTAP \(Data ONTAP 8.2.x and 8.3\)](#)

- Qtree SnapMirror relationships

Volumes with qtrees that are the source of a qtree SnapMirror relationship can be transitioned, but the qtree SnapMirror relationship is not transitioned. A volume with a qtree that is the destination of a qtree SnapMirror relationship can be migrated only after the qtree SnapMirror relationship is broken.

- Disaster recovery vFiler unit

Volumes that are the source of a disaster recovery vFiler unit can be migrated; however, the disaster recovery vFiler unit is not transitioned. A volume that is the destination of a disaster recovery vFiler unit can be migrated only after the disaster recovery relationship is deleted.

- NDMP configuration

After the transition is complete, you must manually set up backup policies for the transitioned volumes in ONTAP.

[Data protection using tape backup](#)

- Synchronous SnapMirror relationships

This feature is not supported in ONTAP; however, the volumes that are part of the relationship can be transitioned.

Considerations for using SnapMirror for transition

You can create data copy schedules and customize the SnapMirror data transfers for transition operations without affecting the existing 7-Mode to 7-Mode SnapMirror or SnapVault operations.

Maximum number of concurrent SnapMirror transfers

During transition, the maximum number of concurrent SnapMirror transfers supported on the 7-Mode and ONTAP systems depend on the number of volume SnapMirror replication operations allowed for a specific storage system model.

For information about the maximum number of concurrent volume SnapMirror transfers for your system model, see the [Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode](#).

Data copy schedules

- The number of concurrent SnapMirror transfers that the tool uses for running the SnapMirror operations (baseline, update, or resynchronization) is based on the schedules you configure while creating the project.
- If different projects are transitioning volumes from the same 7-Mode controller, you must ensure that the data copy schedules do not overlap across different projects.
- You can ensure that your existing backup and disaster recovery (DR) operations are not impacted by the 7-Mode Transition Tool transition operations in the following ways:

- You should create SnapMirror data copy schedules for a project such that it does not overlap with the existing 7-Mode SnapMirror or SnapVault schedules.
- You should configure the number of concurrent SnapMirror transfers to run in such a way that the existing 7-Mode SnapMirror or SnapVault schedules do not fail.

You can also release some transfers by editing the active schedule and modifying the maximum number of concurrent volume SnapMirror transfers to zero.

- You must ensure that the number of concurrent SnapMirror transfers and the throttle configured for the operations (precutover, cutover, and on-demand update) are available on the 7-Mode storage system for the entire duration of the operation.

The cutover operation fails if the final incremental update operation fails even for one of the volumes in the project.

- For secondary projects, after cutover, the incremental SnapMirror updates for the SnapMirror relationship between the 7-Mode primary volumes and the ONTAP secondary volume is based on the 7-Mode to 7-Mode SnapMirror relationship schedule.

You must ensure that there are sufficient concurrent SnapMirror transfers available on the 7-Mode primary controller for these updates to occur.

Using multiple paths for transition

You can specify two paths for transition by using a data copy IP address and a multipath IP address. However, both paths can be used only for load-balancing, not for failover.

Related information

[Considerations for creating a data copy schedule](#)

[Creating a data copy schedule for SnapMirror transfers](#)

Guidelines for deciding when to perform cutover

Because transition cutover is disruptive to clients, you must plan the activity to minimize the downtime. You must schedule the cutover during a low-activity window. You should update the ONTAP volumes and wait for the transfers to complete before disconnecting clients and initiating storage cutover for reducing the downtime.

You must keep monitoring the SnapMirror status for each volume. If the last transfer duration of the previous few updates for the volume is within an acceptable limit, most of the data changes in the volume should have been copied and the time for final data update during cutover should be within the acceptable limit.

You can derive the approximate downtime depending on the number of volumes that are transitioned.

To minimize the cutover time, the network latency between the 7-Mode Transition Tool and storage systems should be minimum. For transitioning a volume SnapMirror relationship, the network latency between the tool and the primary systems should be minimum.

Related information

[Performing on-demand SnapMirror updates](#)

Impact of takeover and giveback on transition

Transition operations, such as transition prepare, start, pause, resume, or complete, fail during a controller takeover or giveback.

If a transition operation fails due to a takeover, you must wait for the giveback to finish, and then run the transition operation again.

If a controller takeover occurs during a baseline transfer, the transfer fails. To resume the baseline transfer from the point where it was aborted, you must wait for the giveback to finish.

Data copy resumes based on the configured schedule.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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