



Preparing for copy-free transition

ONTAP 7-Mode Transition

NetApp
August 29, 2024

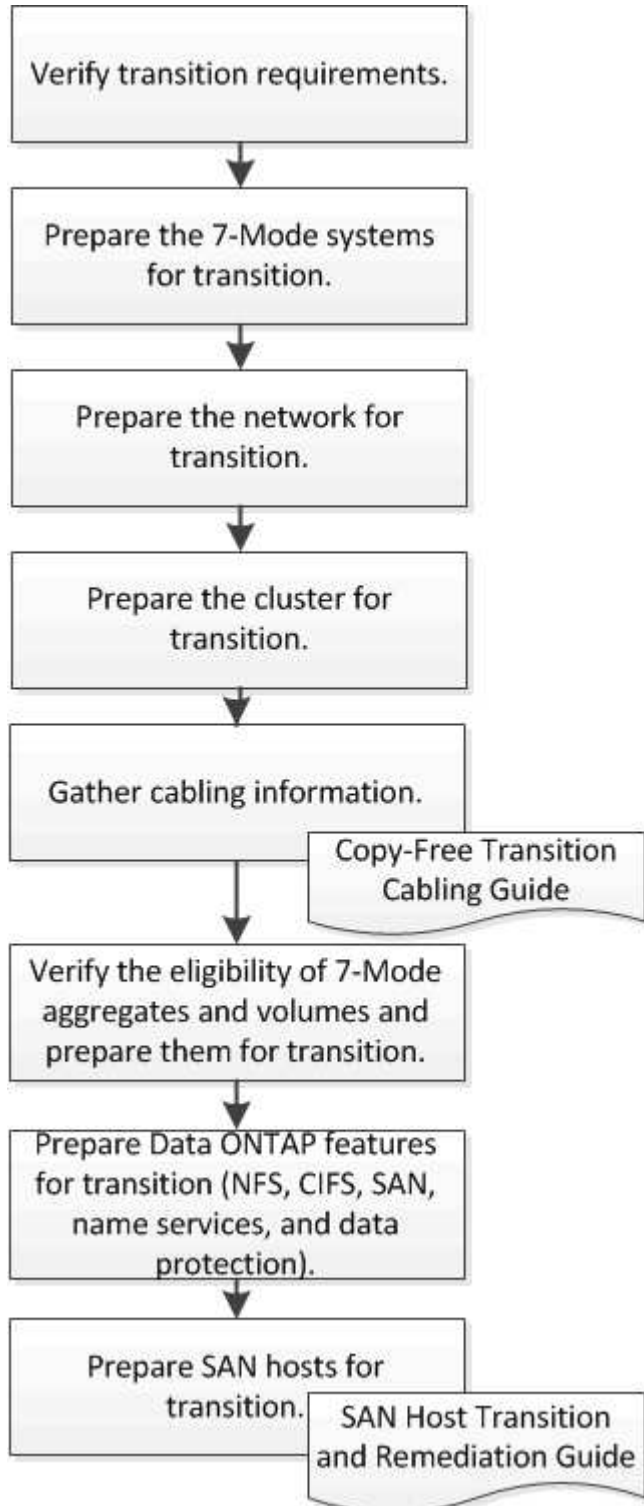
This PDF was generated from https://docs.netapp.com/us-en/ontap-7mode-transition/copy-free/concept_requirements_for_copy_free_transition.html on August 29, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Preparing for copy-free transition 1
 - Requirements for copy-free transition 2
 - Tools and documentation required for copy-free transition 3
 - Port requirements for communicating with the 7-Mode Transition Tool 3
 - Preparing the 7-Mode HA pair for transition 4
 - Setting up the SP or RLM on the 7-Mode systems for copy-free transition 5
 - Preparing the network for transition 8
 - Preparing the cluster for transition 9
 - Gathering cabling information for transition 10
 - Preparing 7-Mode aggregates and volumes for transition 14
 - Preparing to transition name services 20
 - Preparing for NFS transition 23
 - Preparing for SMB/CIFS transition 29
 - Preparing for SAN transition 35
 - Preparing data protection features for transition 39

Preparing for copy-free transition

Before starting the copy-free transition, you must identify the 7-Mode HA pair to transition, understand the requirements and restrictions for migration, and prepare the 7-Mode systems and cluster for transition. You must also be aware of the Data ONTAP features that are supported and unsupported for transition.



Related information

Requirements for copy-free transition

You should be aware of the requirements for 7-Mode systems, clusters, ONTAP releases, and disk shelves for copy-free transition.

Be sure to consult the current 7-Mode Transition Tool *Release Notes* for the latest information about supported target releases and known issues.

7-Mode Transition Tool Release Notes

- **Platform models**

Copy-free transition is supported only on mid-end and high-end FAS systems and IBM N series systems. The [NetApp Interoperability Matrix Tool](#) has the latest information about the supported platforms for 7-Mode systems and the target cluster nodes.

- **Data ONTAP in 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**

7-Mode Transition Tool version 3.3.1 supports transition to the following ONTAP releases by using the copy-free method:

- ONTAP 9.4 and earlier ONTAP 9 releases
- Clustered Data ONTAP 8.3.2 and later 8.x releases **Note:** You cannot use the 7-Mode Transition Tool to transition to ONTAP 9.5 or later using the copy-free method. To do so, you must first transition to ONTAP 9.4 using 7-Mode Transition Tool 3.3.1 and then upgrade your cluster to ONTAP 9.5 or later. 7-Mode Transition Tool 3.3.2 does not support copy-free transitions.

- **HA configuration**

The 7-Mode controllers and target cluster nodes must be in an HA configuration. The HA pairs must be healthy, and none of the nodes can be in takeover mode. Stand-alone controllers are not supported for copy-free transition.

- **Disk shelf models**

The following disk shelf models are supported:

- DS4486
- DS4246
- DS4243



The disk shelf model DS4243 is not supported with ONTAP 9.2 and ONTAP 9.4. This model is supported with all ONTAP 9.2 patch releases starting with ONTAP 9.2P1 and with ONTAP 9.3. 7-Mode Transition Tool 3.3.1 supports transition with the disk shelf model DS4243 for copy-free transition to ONTAP 9.2P1 through ONTAP 9.3.

- DS2246
- DS14mk4 FC (not supported in ONTAP 9.0 and later)
- DS14mk2 AT (not supported in ONTAP 9.0 and later)



The disk shelf model DS14mk2 FC is not supported.

- **Disk firmware**

You must download and install the latest disk qualification package, disk firmware, and disk shelf and ACP firmware on the 7-Mode systems and target cluster nodes.

[NetApp Downloads: Disk Qualification Package](#)

[NetApp Downloads: Disk Drive Firmware](#)

[NetApp Downloads: Disk Shelf Firmware](#)

- **Tool to verify cabling**

After connecting the 7-Mode disk shelves to the target cluster nodes during the transition, you must use Config Advisor to verify the cabling.

[NetApp Downloads: Config Advisor](#)

Tools and documentation required for copy-free transition

The Config Advisor is the required tool for copy-free transition. You should use Config Advisor to verify the cabling of the disk shelves. Additional documentation is also available for SAN host remediation.

Config Advisor

You should use the “Transition” execution profile in Config Advisor to verify the cabling after the 7-Mode disk shelves are connected to the target cluster nodes.

[NetApp Downloads: Config Advisor](#)

Documentation

Describes the pre-transition and post-transition steps that have to be performed on SAN hosts when transitioning using copy-free transition.

[SAN host transition and remediation](#)

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.

Ports that must be open on the cluster

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

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.

Related information

[7-Mode Transition Tool installation and administration](#)

Preparing the 7-Mode HA pair for transition

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

The HA pair must be healthy and none of the nodes must be in the takeover mode, which can be verified by using the `cf status` command. You can also use the NetApp AutoSupport tool to detect any errors or at risk conditions.

1. If HTTPS is not enabled on the storage system, enable HTTPS:

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

HTTPS is enabled by default.

2. 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.

Setting up the SP or RLM on the 7-Mode systems for copy-free transition

If the Service Processor (SP) or the Remote LAN Module (RLM) is not already configured on the 7-Mode storage systems or if you have configured the SP or RLM with an IPv6 address, you must configure the SP or RLM with an IPv4 address.

- SSHv2 must be supported on the host on which the 7-Mode Transition Tool is installed.
- You must have access to the SP or RLM “naroot” account or a Data ONTAP user account with the credentials of the “admin” role or a role with “login-sp” capability.

7-Mode Transition Tool accesses the 7-Mode systems when the systems are halted during transition by using a remote management device that can be the SP or RLM, whichever is available on your system based on the platform model. You must configure the SP or RLM with an IPv4 address. IPv6 configuration is not supported for transition.

Steps

- Configure the SP and provide SP access to the host on which 7-Mode Transition Tool is installed.
 - a. Configure and enable the SP network with an IPv4 address:

```
sp setup
```

```

system1> sp setup
    The Service Processor (SP) provides remote management capabilities
    including console redirection, logging and power control.
    It also extends autosupport by sending
    additional system event alerts. Your autosupport settings are use
    for sending these alerts via email over the SP LAN interface.
Would you like to configure the SP? y
Would you like to enable DHCP on the SP LAN interface? n
Please enter the IP address of the SP []: 192.168.123.98
Please enter the netmask of the SP []: 255.255.255.0
Please enter the IP address for the SP gateway []: 192.168.123.1
Do you want to enable IPv6 on the SP ? n
Verifying mailhost settings for SP use...

```

- b. Verify the SP network configuration settings:

sp status

```

system1> sp status
    Service Processor      Status: Online
    Firmware Version:      1.2
    Mgmt MAC Address:      00:A0:98:01:7D:5B
    Ethernet Link:         up
    Using DHCP:            no
    IPv4 configuration:
    IP Address:            192.168.123.98
    Netmask:               255.255.255.0
    Gateway:               192.168.123.1

```

- c. Provide SP access to the host on which the 7-Mode Transition Tool is installed:

options sp.ssh.access host=7mtt_host

7mtt_host is the host name or IP address of the host on which the 7-Mode Transition Tool is installed.



When you configure the SP, all hosts are granted access by default. You must perform this step if you want to restrict the access to specific hosts.

- d. From the host on which the 7-Mode Transition Tool is installed, log in to the SP:

ssh username@SP_IP_address

When prompted, enter the password for the user name.

The SP prompt is displayed, indicating that you have access to the SP CLI.

- Configure the RLM and provide RLM access to the host on which the 7-Mode Transition Tool is installed.

- a. Configure the RLM network with an IPv4 address:

rlm setup

In the RLM CLI wizard, you must enter the IP address, network mask, and gateway for the RLM.

```
system> rlm setup
    The Remote LAN Module (RLM) provides remote management
capabilities
    including console redirection, logging and power control.
    It also extends autosupport by sending
    additional system event alerts. Your autosupport settings are
used
    for sending these alerts via email over the RLM LAN interface.
Would you like to configure the RLM? y
Would you like to enable DHCP on the RLM LAN interface? n
Please enter the IP address for the RLM []:192.168.123.98
Please enter the netmask for the RLM []:255.255.255.0
Please enter the IP address for the RLM gateway []:192.168.123.1
Do you want to enable IPv6 on the RLM ? n
Verifying mailhost settings for RLM use...
```

- b. Verify that the RLM network configuration is correct:

rlm status

```
system> rlm status
Remote LAN Module      Status: Online
  Part Number:         110-00030
  Revision:            A0
  Serial Number:       123456
  Firmware Version:    4.0
  Mgmt MAC Address:    00:A0:98:01:7D:5B
  Ethernet Link:       up, 100Mb, full duplex, auto-neg complete
  Using DHCP:          no
IPv4 configuration:
  IP Address:          192.168.123.98
  Netmask:              255.255.255.0
  Gateway:              192.168.123.1
```

- c. Provide RLM access to the host on which the 7-Mode Transition Tool is installed:

```
+options rlm.ssh.access host=7mtt_host*
```

7mtt_host is the host name or IP address of the host on which the 7-Mode Transition Tool is installed.



When you configure the RLM, all hosts are granted access by default. You must perform this step if you want to restrict the access to specific hosts.

- d. From the host on which the 7-Mode Transition Tool is installed, log in to the RLM:

```
ssh username@RLM_IP_address
```

When you are prompted, you must enter the password for the user name.

The RLM prompt is displayed, indicating that you have access to the RLM CLI.

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 prepare the cluster to communicate with the 7-Mode Transition Tool and prepare the SVMs for transition. You can transition to a target HA pair that has data aggregates.

- The cluster must already be set up and the target cluster nodes must be joined to the cluster.

Software setup

- The SVMs must be created and assigned to an IPspace.
- You can transition the 7-Mode disk shelves to a target HA pair that has preexisting data aggregates and volumes.

For a two-node cluster, you must have a data aggregate to host the root volumes of the target SVMs. For a cluster with four or more nodes, the root volumes of the SVMs can be hosted either on the target nodes of the transition or on other nodes in the cluster.

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.

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

Gathering cabling information for transition

Before starting copy-free transition, you must gather information about the adapters, ports, disk shelves, and storage connectivity of your 7-Mode controllers, and then plan how to connect the 7-Mode disk shelves to the target cluster nodes.

You must have printed the copy-free transition cabling worksheet.

Copy-free transition cabling worksheet

1. Use Config Advisor to perform a health check on the 7-Mode storage and cabling and collect cabling data.

You should use the 7-Mode Install Checks option from the “Data ONTAP 7 and 8 (7-Mode)” execution profile.

2. Gather the required information about each 7-Mode controller by using the following command:

sysconfig slot_number

You can use the output of this command to identify which ports are used for disk shelf connectivity.

```
host1> sysconfig 3
    slot 3: SAS Host Adapter 3a
              24 Disks:          13440.0GB
              1 shelf with IOM3
    slot 3: SAS Host Adapter 3b
              24 Disks:          13440.0GB
              1 shelf with IOM3
    slot 3: SAS Host Adapter 3c
              24 Disks:          13440.0GB
              1 shelf with IOM3
    slot 3: SAS Host Adapter 3d
              24 Disks:          13440.0GB
              1 shelf with IOM3
```

3. From the cluster, run the following nodeshell command on each node:

system node run -node node_name -command sysconfig -a

You can use the output of this command to obtain information about the available ports and expansion card slots.

4. On the target cluster nodes, plan the ports to be used for connecting the 7-Mode disk shelves:
 - a. Review the available (open) ports.
 - b. Review the expansion card slots.
 - c. Plan the expansion card configuration.

You can plan to move the expansion cards from the 7-Mode systems if they are also supported on the destination platform and ONTAP version. You can also plan for PAM cards, if required.

[NetApp Hardware Universe](#)

- d. Plan the destination ports to use for the disk shelf cabling.

The selection of the destination ports depends on some of the following factors:

- Separate or existing disk shelf stack
- Port availability
- SAS or FC connections
- Availability of on-board ports or expansion cards

5. Go to the data center to physically record the port connections on the 7-Mode controllers and target cluster nodes in the cabling worksheet:
 - a. Record the used ports on the 7-Mode controllers in the cabling worksheet.
 - b. Record the used ports on the target cluster nodes in the cabling worksheet.
 - c. Record the destination ports to be used for connecting the 7-Mode disk shelves, as planned in Step [#STEP_D0CFE719A0384F7FA5D9E73C8EA6C2E7](#).
 - d. Ensure that you have the right cables for connecting the disk shelves.

You should identify any issues with cabling based on the new disk shelf stack location.

- e. Plan for longer cable lengths due to ladder racking or data center requirements.
- f. Label each disk shelf stack and cable on the 7-Mode controllers.

The best practice is to label the 7-Mode disk shelf stacks in case you want to roll back the transition and have to reconnect the disk shelves to the 7-Mode controllers.

Related information

[SAS Disk Shelves Installation and Service Guide for DS4243, DS2246, DS4486, and DS4246](#)

[DiskShelf14mk2 AT Hardware Service Guide](#)

[DS14mk2 FC, and DS14mk4 FC Hardware Service Guide](#)

Copy-free transition cabling worksheet

You can use the copy-free transition cabling worksheet to plan your cabling. You must record information about the ports and disk shelves connected to the 7-Mode controllers and target cluster nodes. You should also record the ports to use for connecting the 7-Mode disk shelves to the target cluster nodes.

7-Mode cabling				Clustered Data ONTAP cabling			
Location: Colorado Floor: Third Rack: 8				Location: Colorado Floor: Fifth Rack: 3			
Module A Ports	Module B Ports	Shelf Type/Asset Tag	Shelf IDs	Module A Ports	Module B Ports	Shelf Type/Asset Tag	Shelf IDs
1a	0a	DS4243/174 263-6	10-13	1a	0a	DS4243/174 233-2	10-11
1b (offline)	0b (offline)	n/a	n/a	1b	0b	DS4243/174 263-6	20-23
1c	0c	DS4243/174 274-9	30-37	1c	0c	DS4243/174 274-9	30-37
1d	0d	DS4243/174 285-6	14-15	1d	0d	DS4243/174 285-6	14-15

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 aggregates and volumes

You must be aware of certain restrictions for transitioning 7-Mode aggregates and 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 \(Data ONTAP 8.2.x and 8.3\)](#)

- SnapLock volumes

The transition of SnapLock volumes is supported for all of the latest ONTAP releases.

- FlexCache volumes

Aggregate and volume states

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

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

FlexClone volumes

The clone hierarchy and storage efficiency are preserved during the copy-free transition. However, you must ensure that the parent FlexVol volume and all of its FlexClone volumes belong to the same vFiler unit. If the FlexClone volumes are in different vFiler units from the parent volume, you must choose one of the following actions:

- Move the FlexClone volumes to the vFiler unit that owns the parent FlexVol volume.
- Split the clones from the parent FlexClone volume, and then transition these volumes as FlexVol volumes.

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 7-Mode systems with 32-bit aggregates

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.

- **32-bit aggregates**
 - [Expanding an aggregate to the 64-bit format](#)
 - [Finding and removing 32-bit volumes and Snapshot copies](#)
- **32-bit volumes or Snapshot copies**

Even if you have only 64-bit aggregates and volumes, some 32-bit or mixed-format FlexVol volumes or Snapshot copies might remain. You must remove these volumes and Snapshot copies before transition.

[Finding and removing 32-bit volumes and Snapshot copies](#)

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 <code>grow-all</code> 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
```

Aggregate space requirements for transition

Before transition, you must ensure that the 7-Mode aggregates have adequate free space. The 7-Mode Transition Tool performs various space checks on the aggregates based on the physical space, logical space, space occupied by Snapshot copies, and space guarantee settings. You must also be aware of the space considerations with Flash Pool aggregates.

Physical space in the aggregates

Transition is blocked if the free space is less than 5% of the physical space in the 7-Mode aggregates. The best practice is to have at least 20% free space in the 7-Mode aggregates before transition.

The additional space is required in the aggregates for the following reasons:

- Creating the aggregate-level Snapshot copy for each 7-Mode aggregate during the export phase
- Testing the workload on the transitioned aggregates with new data in the preproduction testing phase

If you do not have additional space, you can add disks to the 7-Mode systems before transition. If adding disks is not feasible or if you can ensure that only limited amount of data is written on the transitioned volumes during the preproduction phase, the 7-Mode Transition Tool allows you to acknowledge this error and continue with the transition. However, you must continue to monitor the aggregate space during the transition and ensure that the aggregates do not grow in the preproduction testing phase.

Logical space in the aggregates

If the logical space in the 7-Mode aggregates is more than 97% full, 7-Mode Transition Tool throws a blocking error during precheck. You can ignore this error during the planning phase and continue with the transition; however, you must ensure that the logical space used is less than 97% before the export and halt operation by either reducing the size of the volumes in such aggregates or adding more disks to the aggregates. You cannot

ignore this error in the export and halt phase.

Snapshot spill

If the Snapshot copies in the 7-Mode aggregates occupy more space than the allocated space for Snapshot copy reserve, the creation of aggregate-level Snapshot copies in the export and halt operation might fail. 7-Mode Transition Tool throws a blocking error during precheck for this condition. In such conditions, you must delete all the existing aggregate-level Snapshot copies during the planning phase.

If you do not want to delete the existing Snapshot copies, you can ignore this error during the planning phase and continue with the transition; however, you must ensure that the Snapshot copy used capacity percentage is less than 100% before the export and halt operation.

Space guarantee settings

7-Mode Transition Tool throws a blocking error during precheck if the 7-Mode controllers have volumes with the following space guarantee settings:

- Volume-guaranteed volumes with guarantee disabled
- File-guaranteed volumes
- **Volume-guaranteed volumes with guarantee disabled**

In some cases, the space guarantee is disabled for the volume guaranteed volumes because of lack of space in the aggregates.

You must create sufficient free space on the 7-Mode aggregates and then enable space guarantee for such 7-Mode volumes by using the following 7-Mode command:

```
vol options volume_name guarantee volume
```

If you do not want to perform any corrective actions on 7-Mode, you can ignore this error. After the transition, examine the volumes for which guarantee is disabled and enable the guarantee manually by using the following command:

```
volume modify -vserver -volume -space-guarantee volume
```

- **File-guaranteed volumes**

File guarantee is not supported in ONTAP.

If you have file-guaranteed volumes, you must perform one of the following actions:

- If the 7-Mode volumes contain space-reserved LUNs or files, change the space guarantee type of the volumes to volume by using the 7-Mode command:

```
vol options volume_name guarantee volume
```

You must ensure that there is enough free space on the 7-Mode aggregates before running this command.

- If the 7-Mode volumes do not contain any space-reserved LUNs or files, change the space guarantee of the volumes to none by using the following 7-Mode command:

```
vol options volume_name guarantee none If you do not want to perform any corrective actions on 7-Mode, you can ignore this error and continue with the transition.
```

During the transition, if these volumes contain space-reserved LUNs or files, their space guarantee will be

automatically converted to `volume`, but the space guarantee will be disabled initially. You must create sufficient free space on the aggregates and then manually enable the guarantee by using the following command:

```
+ volume modify -vserver -volume -space-guarantee volume
```

+ If the volumes do not contain any space-reserved LUNs or files, their space guarantee will be automatically converted to `none` during the transition.

Additional consideration for Flash Pool aggregates

Transition is not supported if the free space in the SSDs of Flash Pool aggregates is less than 5% of the total disk space of the SSDs. You must either disable the SSD cache or add more SSDs to continue with the transition.

Related information

[Ignorable errors during transition](#)

[Disk and aggregate management](#)

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

[Customizing the transition of 7-Mode configurations](#)

[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 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.

[Customizing the transition of 7-Mode configurations](#)

- 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.

[Customizing the transition of 7-Mode configurations](#)

- 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 all of the vFiler units.
- NFS service must be running on the 7-Mode systems during transition.

Even after client access is disconnected and you prepare to start the export phase, the service must be running on the 7-Mode systems.

- 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.

Related information

[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`
- `waf1.default_nt_user`
- `nfs.mount_rootonly`
- `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`

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).
- Kerberos configuration

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

[Customizing the transition of 7-Mode configurations](#)

[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 disallow access to all NFS clients.

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/mtree1
transition_volume_qtreemount:192.168.35.223:/vol/volstd10/mtree1 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 qtree. However, when the client accesses the qtree through the mount point of its parent volume, the client can write to the qtree 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 qtree qtrees1 when the client accesses the qtree directly or through the mount point of the qtree'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.
- Even after client access is disconnected and you prepare to start the export phase, the CIFS service must be running on the 7-Mode systems.
- The authentication type for CIFS must be Active Directory (AD) or Workgroup.

Cluster prerequisites

- The CIFS license must be added.
- 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.
- The following table identifies which authentication method 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 the CIFS configuration from 7-Mode to ONTAP if the AD domains do not match between the 7-Mode CIFS server and the target SVM CIFS server.

The tool triggers an ignorable blocking error when an AD domain name mismatch is detected. To proceed with the transition, you can acknowledge the blocking error.

- The CIFS server must be manually configured before the apply configuration (precutover) phase.

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

If you want to...	Do the following...
Transfer or preserve the CIFS server identity to the target SVM	<div data-bbox="844 159 1484 338" style="border: 1px solid #ccc; padding: 10px; margin-bottom: 10px;"> <p>You have the following two options to create the CIFS server:</p> </div> <p>a. Applicable for all versions of ONTAP:</p> <ul style="list-style-type: none"> ◦ Before the SVM provision phase, 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 “SVM Provision” and “Export & Halt” phases with the new temporary identity. This action is required to read CIFS configurations from 7-Mode during the SVM Provision and “Export & Halt” phases.</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 “SVM Provision” operation, and then perform the “Export & Halt” operation to enable client access to ONTAP volumes. <p>b. Applicable for ONTAP releases 9.0 through 9.5:</p> <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 “SVM Provision” operation.</p> <ul style="list-style-type: none"> ◦ After the “import” phase, you can run the <code>vserver cifs modify</code> command on the target cluster to replace the target SVM CIFS identity with the 7-Mode CIFS identity.

If you want to...	Do the following...
Use a new identity	<ul style="list-style-type: none"> • Before the “SVM Provision” phase, 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 “SVM Provision” and “Export & Halt” phases. <p>This action is required to read CIFS configurations from 7-Mode during the “SVM Provision” and “Export & Halt”.</p> <ul style="list-style-type: none"> • After verifying these conditions, you can perform the “SVM Provision” operation. <p>You can then test the SVM configurations, and then plan to perform the storage cutover.</p>

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.

Related information

[Customizing the transition of 7-Mode configurations](#)

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.

- 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.

- You must ensure that the number of local users and groups in 7-Mode does not exceed the local users and groups limit for ONTAP.

You must contact technical support if the number of local users and groups in 7-Mode exceeds the limit defined in ONTAP.

- 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 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.

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](#)

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 FC zone planner must be generated by using the Collect and Access feature of the 7-Mode Transition Tool
- The FC zone script file must be accessible.
 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.

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.

2. Perform the pretransition steps on the host.

SAN host transition and remediation


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.

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 the hosts are quiesced. If the volume does not have sufficient free space available, the required amount of space must be added to the 7-Mode volume.

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

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 space that has to be added is the size of the Snapshot reserve

for the volume. After space is added to the volume, you can use the `lun transition start` command to transition the LUNs.

Related information

[Recovering from a failed LUN transition](#)

[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.

Preparing the cluster for transitioning volume SnapMirror relationships

For transitioning 7-Mode volume SnapMirror relationships, you must add the SnapMirror license to the source and destination clusters. You must also create a cluster peer relationship between the clusters to which the primary and secondary volumes of the SnapMirror relationships are transitioned and create the SnapMirror schedules.

You must have recorded the SnapMirror schedules defined in the `/etc/snapmirror.conf` file of the 7-Mode secondary system.

Steps

1. Add the SnapMirror license on both the source and destination clusters:

```
system license add license_code
```

2. From each cluster, create the cluster peer relationship.

[Cluster and SVM peering express configuration](#)

3. Create schedules on the secondary SVMs that match the schedules in the 7-Mode secondary system:

```
job schedule cron create
```

Related information

[ONTAP 9 commands](#)

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 perform a copy-free transition of primary and secondary HA pairs. You must then manually set up the volume SnapMirror relationships after transition.

[Transitioning a SnapMirror relationship](#)

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.

Related information

[Customizing the transition of 7-Mode configurations](#)

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.