



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_copy_based_transition_using_7mtt.html on February 11, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Copy-Based Transition Guide 1
 - Transition overview 1
 - Collecting and assessing the inventory information 4
 - Copy-based transition workflow 14
 - Data and configuration migration process 15
 - Preparing for copy-based transition 28
 - Migrating data and configuration from 7-Mode volumes 70
 - Transitioning volumes by using the 7-Mode Transition Tool CLI 89
 - Performing manual post-transition tasks 111
 - Managing a transition project 117
 - Troubleshooting issues 129

Copy-Based Transition Guide

This guide describes how to assess 7-Mode controllers, hosts, and applications for transition and perform a copy-based migration of data and configuration from 7-Mode systems to ONTAP by using the 7-Mode Transition Tool.

Transition overview

Transitioning to clustered ONTAP involves identifying your current environment, defining the transition scope, designing the optimal configuration of the destination systems, planning how to migrate data and configurations, and making necessary environmental updates.

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](#)

You must first collect information about the current environment, including both the current storage environment as well as the hosts and applications. After collecting the storage inventory, you can assess the current features and functionality that are used and identify any differences in the ONTAP version selected. You can use 7-Mode Transition Tool to perform these tasks.

You can then define your migration project. This involves identifying what volumes and workloads you want to migrate together, then you can design your cluster, and plan your transition. You can begin planning your migration by first selecting the optimal migration method. When designing your cluster, you can use the information gathered about your storage environment to map existing configuration items to their equivalent items in ONTAP. For example, a 7-Mode volume should be mapped to an SVM and a volume in ONTAP and an IP address that will be transitioned should be mapped to LIFs. You should also determine if any environmental changes have to be made on the ONTAP system.

Implementation includes deploying and configuring the cluster, migrating data, applying configuration changes, disconnecting clients and reconnecting to the destination ONTAP system, verifying access, as well as performing any necessary environmental changes.

Related information

[NetApp Documentation: ONTAP 9](#)

Copy-based transition by using the 7-Mode Transition Tool

The 7-Mode Transition Tool enables you to perform copy-based transition by collecting information and assessing 7-Mode controllers, hosts, switches, and applications for transition. You can then migrate your data and configurations from 7-Mode to ONTAP.

The 7-Mode Transition Tool runs on a Windows or Linux system. The 7-Mode Transition Tool provides both a web interface and a command-line interface for managing your transition operations.

Collecting and assessing ONTAP systems, hosts, switches, and applications

You can perform the following collect and assess tasks by using the 7-Mode Transition Tool:

- Collect inventory information from ONTAP systems (7-Mode controllers and nodes in the cluster), hosts, switches, and host applications.
- Generate the FC Zone plan to configure the zones for grouping the initiator hosts and targets.
- Assess the features and functionalities of the 7-Mode systems, and identify how these features and functionalities work in the ONTAP version selected for transition.

Moving data and configurations from 7-Mode to ONTAP

Copy-based transition uses SnapMirror technology to copy 7-Mode volumes and configurations from Data ONTAP 7G and 7-Mode to ONTAP.

You can perform the following tasks by using the 7-Mode Transition Tool for copy-based migration:

- Migrate a group of 7-Mode stand-alone volumes or a group of 7-Mode volumes in volume SnapMirror relationships from systems running different 7-Mode versions to any version from ONTAP 8.2.x and later supported releases.
- Run prechecks on volumes included in a transition project to verify their compatibility for transition and view possible corrective actions.
- Apply 7-Mode configurations to ONTAP before disconnecting client access, reducing the downtime.

Copy-based transition supports the transition of NAS and SAN configurations.



SAN transition is supported only to ONTAP 8.3 and later supported releases.

Related information

[Requirements for copy-based transition](#)

Transition terminology

Understanding the transition terminology related to the 7-Mode Transition Tool helps you to understand the transition process.

• Collect and assess

Before transitioning data and configurations to ONTAP, you must collect information about the storage environment that includes storage systems, hosts, and applications. You must then assess the features and functionalities of these systems, and identify how these features and functionalities work in the ONTAP version selected for transition.

• Migrate

Refers to transitioning data and configurations from the 7-Mode volumes to ONTAP. Migration should be performed after assessing the 7-Mode controllers.

• Project

In the 7-Mode Transition Tool, a project enables you to configure and manage the transition of a group of volumes.

Transition projects are of three types: stand-alone, primary, and secondary.

- **Stand-alone project**

Volumes in this project are either not in any SnapMirror relationship with volumes in other storage systems, or you want to transition these volumes without retaining their SnapMirror relationships.

- **Primary project**

Volumes in this project are the source volumes of a SnapMirror relationship.

- **Secondary project**

Volumes in this project are the destination volumes of a SnapMirror relationship.

- **Project group**

In the 7-Mode Transition Tool, a project group is a logical container that you can use to hold related migration projects. There is always one default group with Default_Group name that exists in the system.

- **Transition peer relationship**

Transition peer relationship is an authorization mechanism that enables SnapMirror to establish relationships between a 7-Mode storage system and an SVM in the cluster for copying data from the 7-Mode volumes to the ONTAP volumes for transition. A transition peer relationship can be created only by a cluster administrator and is initiated from the cluster.

- **Transition data protection relationship**

Transition data protection (TDP) relationships are volume SnapMirror relationships that are established between a 7-Mode system and an ONTAP system.

Transition data protection is supported only for transition. You should not keep this relationship for long periods of time because it is not possible to resynchronize data from an ONTAP volume back to a 7-Mode volume.

Limits for transition

When you transition volumes using the 7-Mode Transition Tool, you should consider certain limits for assessment, such as the number of controllers and hosts that can be assessed simultaneously. For migration, you should review the number of volumes in a project and number of projects that can be run simultaneously.


The following table lists the number of controllers and hosts that can be assessed simultaneously:

| Parameter | Maximum limit |
|-----------------------|---------------|
| Number of controllers | 50 |
| Number of hosts | 20 |
| Number of ESXi hosts | 10 |



The maximum number of 7-Mode controllers in a single assessment depends on the number of objects, such as number of volumes, qtrees, quota, and exports.

The following table lists the number of volumes in a project and the number of projects that can be run simultaneously during migration:

| Parameter | Maximum limit |
|---|--|
| Number of volumes in a project | 160 |
| Recommended limit for the total number of volumes across all active projects in the web interface at any given time | 240 |
| Maximum number of active project at any given time | 50 <div> If you have 50 active projects, you can replicate data from 50 7-Mode storage systems and their associated SnapMirror destinations.</div> |
| Number of data copy schedules per project | 7 |


Collecting and assessing the inventory information

You can collect inventory information from controllers, hosts, and FC switches. You can then assess features and functionalities of these systems, and identify how these features and functionalities work in the ONTAP version that is selected for transition.

You can collect inventory information in two ways:

- If your environment security allows it, you can install the 7-Mode Transition Tool, and then use it to collect the inventory information.
- You can import the inventory XML report generated by the Inventory Collect Tool, and then perform the assessment.

In both cases, you must use Inventory Collect Tool 3.3 to collect the inventory.



Storage, host, and FC switches version requirements for transition assessment

You must be aware of the versions of Data ONTAP operating in 7-Mode, hosts, and FC switches that are supported 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.

[NetApp Interoperability Matrix Tool](#)

Preparing the 7-Mode systems and hosts for transition assessment

You must ensure that the 7-Mode systems and hosts meet certain network and protocol requirements for successfully generating an assessment report.

Steps

1. Enable HTTPS on the 7-Mode system:

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

2. Enable TLS on the 7-Mode system:

```
options tls.enable on
```



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

3. Enable SSL and disable SSLv2 and SSLv3 on the 7-Mode system:

- a. Set up and start SSL:

```
secureadmin setup ssl
```

- b. Enable SSL:

```
options ssl.enable on
```

- c. Disable SSLv2 and SSLv3:

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

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



The best practice is to disable SSLv2 and SSLv3 to avoid security vulnerabilities.

4. Enable SSH on the 7-Mode system:

- a. Set up SSH on the 7-Mode system:

```
secureadmin setup -f ssh
```

The `-f` option forces the setup to run even if the SSH server is already configured.

- b. Enable SSH:

```
secureadmin enable ssh2
```

- c. Enable password authentication on the SSH server:

```
options ssh.passwd_auth.enable
```

- d. Enable SSH access to the host: + **options ssh.access**

5. Prepare your Windows host systems:

- Enable WMI access.

For more information about enabling WMI access, see the host documentation.

- If you have Windows Server 2003, verify that you have installed the Microsoft Fibre Channel Information Tool (fcinfo) package and run the tool once on your Windows host system.

This tool enables you to collect the HBA configuration information of the host.

- If the system on which the 7-Mode Transition Tool is run does not belong to a domain, verify the following:
 - The host system belongs to a domain.
 - If the host has a local user, and the user name for that user is in the following format:

`SystemName\Username`

6. Enable SSH on the Linux or ESXi host.

For more information about enabling SSH, see the host documentation.

7. Verify that you have installed the latest NetApp Host Utilities software for each host.

For information about downloading and installing the NetApp Host Utilities software, see the NetApp Support Site.

8. Verify that all the hosts and storage systems can be reached by the Windows system from which the 7-Mode Transition Tool is run.

Assessing controllers and hosts

You can collect and assess information about the controllers and hosts by using either the 7-Mode Transition Tool or the Inventory Collect Tool, depending on the security regulations in your environment.

- The 7-Mode Transition Tool collects inventory information about controller and hosts by adding the systems or by using the inventory report generated by the Inventory Collect Tool.

The 7-Mode Transition Tool then assesses inventory information and creates the transition assessment report.

- You must be aware of the following considerations when performing transition assessment:
 - You should not perform both assessment and migration operations simultaneously on a controller.
 - You should avoid performing assessment operations on active storage controllers during peak hours.

Generating an assessment report by adding systems to the 7-Mode Transition Tool

You can collect inventory information for controllers, hosts, and FC switches by adding the systems to the 7-Mode Transition Tool. You can then create an assessment report to assess the features and functionalities of those systems, and to identify how they work in the ONTAP version selected for transition.

- The user name for the storage system and hosts must have sufficient privileges to execute the commands listed in the readme file.

The readme file is located at `7-Mode_Transition_Tool_installed_location\bin\ict`.

- You must have prepared the 7-Mode systems, hosts, and FC switches for transition assessment.
- For assessing Windows systems, you must have a domain user account.
- If you are adding multiple systems for assessment, you must create a text file that is encoded in the ASCII or UTF-8 format and must contain the system details in the form of one system per line.

Each system details must be in the following format:

```
(ontap|windows|vmware|linux|cisco|brocade) ://[ (user|domain_user) [:password]@] (host_name|ip)
```

- The controller or host must be accessible by the system on which the 7-Mode Transition Tool is installed and run.
- All features must be configured or their license enabled so that the workbook can contain inventory information about these features.
- The user name for the storage system must have sufficient administrative privileges to collect inventory information.
- All host names and storage system configurations, such as CIFS share names, user names, and group names, must be in the UTF-8 format.

If the 7-Mode Transition Tool service or the system on which this tool is installed is restarted, then the system details added to the tool are lost and the system must be added to the tool again.

Steps

1. If you want to use the latest Interoperability Matrix (IMT) data for transition assessment:
 - a. Download the IMT data from the Interoperability Matrix, and then save it:
 - i. From the Reports menu, click **Complete Daily Exports**.
 - ii. In the Complete Daily Exports dialog box, enter FAS in the search field.
 - iii. Download the ONTAP SAN Host excel file, and then save it. [NetApp Interoperability Matrix Tool](#)
 - b. From the CLI, import the IMT data by using the `transition imt import` command.
 - c. Verify that the import is successful by using the `transition imt show` command.

Troubleshooting: If the IMT data import operation fails, you can revert to the previous data by using the `transition imt restore` command.

2. Log in to the 7-Mode Transition Tool, and then click **Get Started** in the Collect and Assess section.
3. Click **Add Systems**.
4. In the Add System window, perform one of the following actions:
 - Add a single system:
 - i. Enter the fully qualified domain name (FQDN) or IP address of the system.

ii. Enter the user name and password for the specified system.

iii. Select the system type:

- Data ONTAP storage systems
- Hosts: Microsoft Windows, Red Hat Linux Enterprise, and VMware ESXi
- FC switches: Cisco and Brocade

- Add multiple systems by clicking **Browse**, and then selecting the text file that contains the credentials for multiple systems.

5. Click **Add**.

If the assessment status of a system is Ready, then you can perform transition assessment for that system.

6. Generate the transition assessment report:

- a. Select the systems for transition assessment.
- b. Click **Create Transition Assessment Report**.
- c. In the Create Transition Assessment Report dialog box, select the Data ONTAP version of the target cluster.
- d. Specify a prefix for the file name of the reports.
- e. Click **Generate Report**.

The assessment workbook (report name appended with "AssessmentWorkbook") and assessment executive summary (report name appended with "AssessmentExecutiveSummary") reports are generated in XML format.

+ You can access the assessment workbook, assessment summary, and inventory XML files that are used to generate the assessment report from the ...etc/webapp/transition-gui/tmc folder.

7. View the assessment workbook in Microsoft Excel and assessment executive summary in Microsoft Word by using Microsoft Office 2007 or later versions.

In the assessment workbook, see the Transition Feasibility (CBT), Config Precheck Summary, Config Precheck Details, and CBT Precheck Summary tabs for copy-based transition assessment details.

In the assessment executive summary, see the Copy-Based Transition Feasibility section for controller-level assessment details.

You might have to enable macros in Excel to view the assessment workbook.

In the data collection summary of the assessment workbook, if the access status of a system is `FAILED`, then the inventory information for that system is invalid. In the assessment executive summary, the value of some of the fields of this system is displayed as `Not Assessed`.

Related information

[Downloading transition log files](#)

Generating an assessment report by importing the inventory report XML

You can import the inventory XML report generated by the Inventory Collect Tool to assess the features and functionalities of hosts and controllers. You can then identify how

these host and controllers work in the ONTAP version selected for transition by creating an assessment report.

- You must have run the Inventory Collect Tool and generated the inventory report XML file.



To assess hosts and controllers for copy-based transition, you must use Inventory Collect Tool 3.3 to collect inventory.

- You must have prepared the 7-Mode systems and hosts for transition assessment.

The systems that you want to assess need not be reachable while importing the inventory report and performing transition assessment.

Steps

1. Log in to the 7-Mode Transition Tool, and then click **Collect and Assess** in the home page.
2. Click **Import Inventory Report XML**.
3. Click **Browse**, and then select the XML report generated by the Inventory Collect Tool.
4. Click **Import**.

The assessment status of the system shows `Imported;Ready`.

5. Select the system for which you want to perform transition assessment.
6. Click **Create Transition Assessment Report**.
7. In the Create Transition Assessment Report dialog box, select the Data ONTAP version of the target cluster.
8. Specify a prefix for the file name of the reports.
9. Click **Generate Report**.

The AssessmentWorkbook and AssessmentExecutiveSummary reports are generated in XML format.

10. View the AssessmentWorkbook report in Microsoft Excel and the AssessmentExecutiveSummary report in Microsoft Word by using Microsoft Office 2007 or later versions.

To view the AssessmentWorkbook report in Microsoft Excel, you might have to enable macros in Excel.

Related information

[Downloading transition log files](#)

Generating an FC zone plan

For FC switches, you must generate an FC zone plan as part of the transition assessment report to configure the zones for grouping the initiator hosts and targets after the migration.

- The 7-Mode system, hosts, and the cluster must be connected to the same switch.

[Supported configurations for generating an FC zone plan](#)

- You must have created the required target SVMs and FC LIFs on the cluster.
- The FC LIFs that are created on the target SVMs will have WWPN, which are different from 7-Mode WWPNs. Therefore, perform FC zoning while you transition SAN for FCP.

Steps

1. In the Collect and Assess section, click **Add Systems**.
2. In the Add System window, perform one of the following actions:
 - Add a single system:
 - i. Enter the FQDN or IP address of the system.
 - ii. Enter the user name and password for the specified system.
 - iii. Select the system type:
 - Data ONTAP storage systems
 - Hosts: Microsoft Windows, Red Hat Linux Enterprise, and VMware ESXi
 - FC switches: Cisco and Brocade
 - Add multiple systems by clicking **Browse**, and then selecting the text file that contains the credentials for multiple systems.
3. Click **Add**.

If the assessment status of a system is Ready, then you can perform transition assessment for that system.

4. Generate the transition assessment report with the FC zone plan:
 - a. Select the systems, including the required FC switches, for transition assessment.
 - b. Click **Create Transition Assessment Report**.
 - c. In the pair systems and start FC zone planner dialog box, select the 7-Mode systems (single controller or HA pair), the switch connected to the cluster, and an SVM in the cluster.
 - d. Click **FC Zoning for paired systems**.
 - e. In the Create Transition Assessment Report dialog box, select the Data ONTAP version of the target cluster.
 - f. Specify a prefix for the file name of the reports.
 - g. Click **Generate Report**.

The FC zone plan is generated as a `.zip` file. The plan contains zones created per the igroup configurations on the 7-Mode systems. Each zone contains a single initiator WWPN and multiple SVM target WWPNs.

You must use the FC zone plan for configuring the zones to group the initiator hosts and targets for providing data access from the cluster.

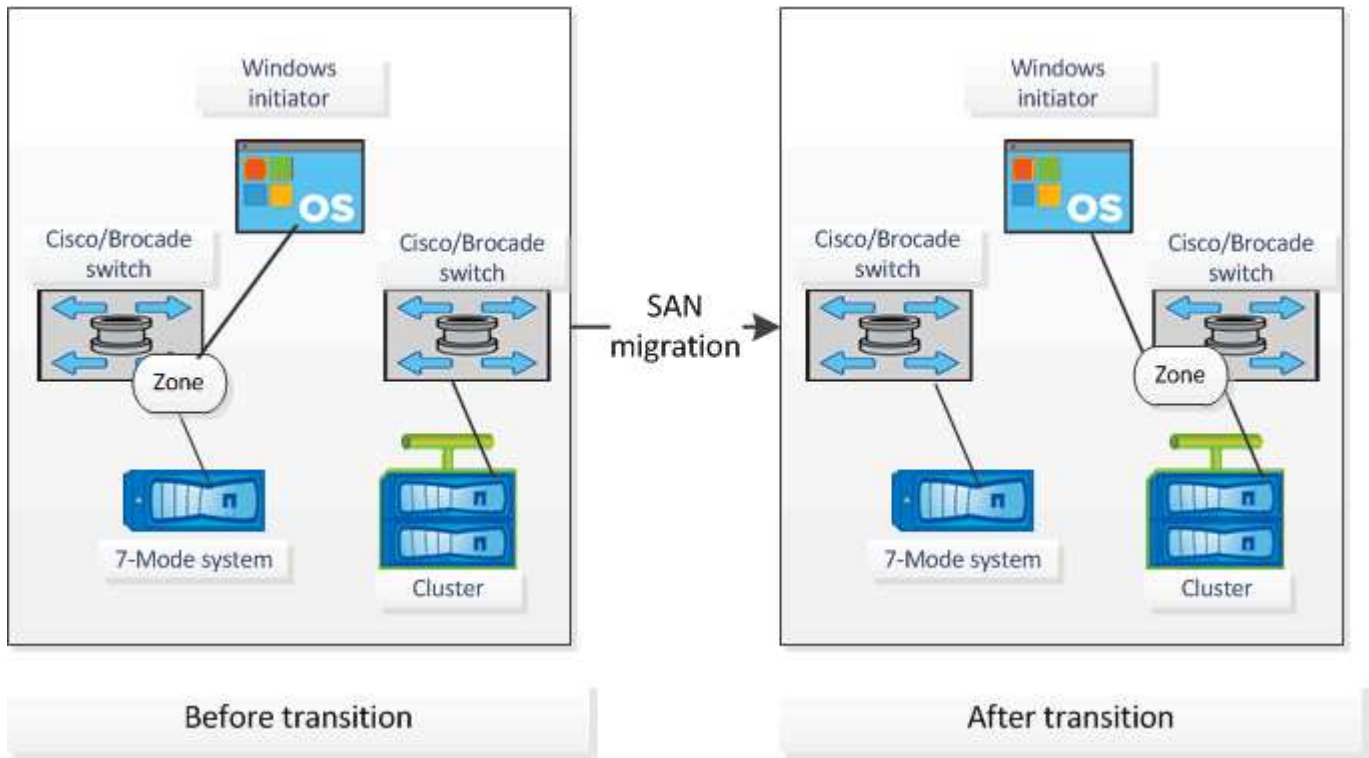
Supported configurations for generating an FC zone plan

You must be aware of the supported configurations of 7-Mode systems, hosts, FC switches, and the cluster to generate the FC zone plan. You should use the plan to configure zones for the cluster after migration.

The 7-Mode systems (single controller or an HA pair), hosts, and cluster can be connected either to the switches in the same fabric or different fabrics, depending on the data center requirements.

The following figure illustrates a configuration in which the 7-Mode systems, hosts, and cluster are connected to the switches in the same fabric:

The following figure illustrates a configuration in which the 7-Mode systems and cluster are connected to switches in different fabrics:



How you can use the assessment executive summary for transition assessment

The transition executive summary provides a summary of the 7-Mode controllers, hosts, and FC switches in your environment. It provides an assessment report of the current features and functionality that are used and recommends the transition methodology for each volume within your storage environment. You can use the summary to plan your transition.

The executive summary has the following main sections:

Target cluster

This section lists the ONTAP version of the target cluster that you selected during assessment.

Data collection summary

You can view the list of 7-Mode controllers, hosts, and switches for which you have collected information. You can view the ONTAP version and model details of the 7-Mode controller. You can also view the OS type, version, and model of the hosts.

Transition feasibility and recommended transition methodology

This section provides a summary of the prechecks that are run on each controller and the feasibility of

transition at the controller and volume level. The volumes that belong to vFiler units that are in the `stopped` or `inconsistent` state or the volumes that are offline or restricted are not included for assessment. The report displays the count of errors and warnings reported in the precheck against each controller. You should review these errors and warnings and resolve any issues before transitioning. Details of these prechecks are available in the Config Precheck Summary tab of the assessment workbook.

Based on the volume and controller configurations and the precheck summary, the executive summary provides a recommendation about the best transition methodology for each assessed volume. For example, you cannot transition 7-Mode traditional volumes or FlexCache volumes because these features are not supported in ONTAP.

For most of the configurations, the 7-Mode Transition Tool is the recommended tool for transition. However, there are some workloads that cannot be transitioned by using the 7-Mode Transition Tool, and for those you should use an application-based or host-based migration method.

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

Storage inventory

This section provides the following information:

- **Storage objects:** Provides information about the number of storage objects, such as volumes, qtrees, LUNs, vFiler units, SnapMirror relationships, shares, and exports, in each controller.
- **Storage utilization:** Provides information about the used space, available space, and space utilized by the 7-Mode controllers.
- **Licenses:** Provides the list of feature licenses enabled on each controller.
- **Protocol configuration:** Provides details about the protocols configured on the controllers, such as CIFS, NFS, and SAN protocols, and the versions.
- **SnapMirror interconnectivity:** Provides information about the controllers or volumes that are either the source or destination of a SnapMirror relationship.

You can use this information to identify controllers that are in SnapMirror relationships with the controllers listed in the report, but are not included for assessment.

- **SnapVault interconnectivity:** Provides information about the controllers, volumes, or qtrees that are either the source or destination of a SnapVault relationship with the specified controller, volumes, or qtrees in the controller.

SVM consolidation considerations

If you are considering consolidating volumes from different vFiler units or 7-Mode controllers to a single SVM, you can use the information provided in this section to identify whether 7-Mode systems are bound to different Active Directory domains, have different NIS configurations, have the same volume names or CIFS share names, and how many local users and groups there are in each controller. These considerations are important in identifying issues when consolidating different vFiler units or 7-Mode controllers and planning the transition to a single SVM.


Data collection errors

This section provides details about controller and host information that could not be collected by the 7-Mode Transition Tool and the reason for the failure. Details of the data collection errors are available in the Data Collection Errors tab of the assessment workbook. You can resolve these errors and assess the systems

again.

Copy-based transition workflow


The copy-based transition workflow involves preparing for migration, migrating data and configurations, and performing post-transition configurations.



Data and configuration migration process

The data and configuration migration process using the 7-Mode Transition Tool consists of the following phases: preparation, baseline data copy, apply configuration (precutover), and storage cutover. If you have SnapLock volumes for Chain of Custody verification, Chain of Custody verification is additional phase after the cutover.

The following image shows the different phases in the migration process:



Preparation

In this phase, prechecks are run to verify feature functionality. The process checks the 7-Mode storage systems to verify that the volumes and configuration are ready to be migrated to ONTAP. It checks that the cluster is configured properly and can support the transition. Any errors must be resolved before continuing with the transition. Although the tool allows you to continue without resolving warnings, you must understand the impact of the warnings before proceeding with the transition. You can run the prechecks multiple times to verify that all of the errors have been resolved.

Although the precheck step and assessment steps performed during the assessment appear to be similar, there are differences. The precheck step is a more detailed test that is focused on the specific storage systems that have been identified as the migration source (7-Mode) and destination (ONTAP) systems. The assessment step only evaluates the migration source systems, checking feature and functionality differences with ONTAP.

Baseline data copy

New volumes are created on the SVM, a SnapMirror relationship is established between the 7-Mode and ONTAP volumes, and a baseline transfer is performed. After the baseline is complete, incremental transfers are automatically run according to a user-defined data copy schedule. Clients and servers accessing the source storage remain online while this step is completed.

Copying data requires CPU, memory, and storage access, which results in additional resources being used on the source storage system. It is a best practice to schedule data copy activity to occur during off-peak times (preferably, CPU usage should be around 50%).

Apply configuration (precutover)

This phase includes SnapMirror incremental transfers; configuration information is applied to the ONTAP system, SVM, and volumes. Optionally, you can also test the ONTAP volumes that are being transitioned

before storage cutover.

Although a majority of the configuration is applied, some actions are deferred to storage cutover: for example, applying quotas.

The 7-Mode IP addresses selected for the transition are created in the administrative down state. The new IP addresses selected for the transition are created in the administrative up state. These new IP addresses can be used to verify data access during precutover testing.

It is a best practice to run the apply configuration (precutover) phase a few days or weeks before the planned cutover window. This activity helps to verify that all of the configurations are applied properly and whether any changes are required.



Although incremental updates are not required, it is a best practice to perform an incremental transfer as close to the storage cutover as possible to minimize the time that clients are disconnected.

Storage cutover

At a high level during storage cutover, clients are disconnected, a final data transfer is performed, the SnapMirror relationship is broken, and clients are manually reconnected.

Disconnecting clients or servers from the source storage volume prevents additional writes from being performed while the final copy is being executed. Before disconnecting clients, it is a best practice to perform an incremental update to minimize the downtime.

Storage access must be disconnected only for the volumes that are being migrated. Access to storage can be discontinued from the storage side or the client side. The best practice is to discontinue connectivity from the storage side. For example, if a CIFS client is accessing a volume named “user01” on a 7-Mode storage system, you can use the `cifs terminate -v user01` command to disable access to all of the CIFS shares on the volume (discontinuing client access from the storage side). The IP addresses, mount points, or even share names might change as a result of the migration, and therefore client access might be discontinued from the client side as well. As long as clients cannot write any new data to the storage container that is being migrated, you can use either or both of these methods for discontinuing access.

After clients are disconnected, the 7-Mode Transition Tool executes a final copy so that both the source and destination datasets are at parity. The 7-Mode Transition Tool configures the data LIFs on the SVM. Some configuration changes that were not transitioned during precutover, such as applying SAN configurations and quotas, are also applied to the SVM at this time.

After storage cutover is complete, you can manually reconnect the clients and validate data access. Validating data access involves verifying that clients are accessing the ONTAP system properly and that all permissions are working as expected.

Chain of Custody verification for SnapLock volumes

You can trigger the Chain of Custody operation for the SnapLock volumes in the project after the transition is complete. This operation is not mandatory and is required only if Chain of Custody verification is essential for the transition of SnapLock volumes. You can perform this operation for all SnapLock volumes in the project or for a subset of SnapLock volumes in the project. The Chain of Custody verification is supported for both compliance and enterprise SnapLock volumes. The Chain of Custody verification is supported only for read-write SnapLock volumes, and is not supported for read-only SnapLock volumes.



The Chain of Custody verification is not supported for SnapLock volume that have file names with non-ASCII characters.

The verification workflow is supported only in the 7-Mode Transition Tool GUI and is not supported in the CLI workflow.

The Chain of Custody verification operation performs the following:

- Enumerates all of the WORM files from 7-Mode volumes
- Calculates the fingerprint for each WORM file enumerated previously on both 7-Mode volumes and transitioned ONTAP volumes
- Generates a report with details about the number of files with matched and unmatched fingerprints, and the reason for the mismatch

Fingerprint data for all WORM files is stored in an ONTAP volume provided during the planning phase.



Based on the number of files on the 7-Mode volumes, the Chain of Custody verification process can take a significant amount of time (days or weeks).

How you transition a stand-alone volume

Transitioning a stand-alone volume includes different phases: preparation, data copy, apply configuration (precutover), and storage cutover. After completing transition, you must perform some post-transition steps before resuming client access. Understanding what occurs during each phase helps you manage your transition efficiently.

| Phase | Steps |
|-------------|---|
| Preparation | <ol style="list-style-type: none">1. Gathering information2. Performing the precheck3. Creating data copy schedules |
| Data copy | <ol style="list-style-type: none">1. Creating the ONTAP volumes as read-only2. Creating a transition peer relationship3. Establishing a SnapMirror relationship4. Performing a baseline transfer5. Performing scheduled incremental updates |

| Phase | Steps |
|--|---|
| Precutover | <ol style="list-style-type: none"> 1. Breaking the SnapMirror relationship 2. Applying configurations to the SVM 3. Configuring data LIFs on the SVM 4. Testing data and configurations (manual and only for precutover RW) 5. Resynchronizing ONTAP volumes with corresponding 7-Mode volumes |
| Storage cutover | <ol style="list-style-type: none"> 1. Disconnecting client access (manual) 2. Performing a final SnapMirror update 3. Breaking the SnapMirror relationship 4. Removing 7-Mode IP addresses and setting the data LIFs to the up state on the SVM 5. Taking the source volume offline <p>After cutover, performing post-transition steps and enabling client access (manual)</p> |
| Chain of Custody verification for SnapLock volumes | <ol style="list-style-type: none"> 1. Enumerating all of the WORM files from 7-Mode volumes 2. Calculating the fingerprint for each WORM file on the 7-Mode volumes (enumerated in the previous step) and calculating the fingerprint for the corresponding WORM file on the transitioned ONTAP volumes 3. Generating a report with details about the number of files with matched and unmatched fingerprints, and the reason for the mismatch |

Preparation phase

In this phase, information about the 7-Mode system and the cluster, volumes, and IP addresses is collected. The 7-Mode Transition Tool performs the following tasks in this phase:

1. Collects and adds 7-Mode storage system and volume information.
2. Runs the transition precheck.
3. Collects and adds cluster, SVM, and aggregate information.
4. Collects IP addresses that must be configured on the SVM:
 - Selects the IP addresses that exist on the 7-Mode system.
 - Specifies new IP addresses that must be configured on the SVM. NOTE: Transitioning of iSCSI and FC LIFs (SAN) is not supported by the tool. You must manually configure SAN LIFs on the SVM before transition.
5. Creates data copy schedules for baseline copy and incremental updates.

6. If the project contains SnapLock volumes, collects information about the read-write SnapLock volumes for which Chain of Custody verification is required and the details of the ONTAP volume that stores the fingerprint data that is generated during the Chain of Custody verification operation.



The Chain of Custody verification operation is supported only for volumes with file names that have only ASCII characters.

7. Plans configuration transition by selecting the 7-Mode configurations that must be transitioned to the target SVM and target volumes.

You should not modify the objects (volumes, IP addresses, system information, and so on) on the controller after fixing the errors and warnings that are reported during the precheck.

Data copy phase

In this phase, data from the 7-Mode volumes is copied to the ONTAP volumes. The 7-Mode Transition Tool performs the following tasks in this phase:

1. Creates the ONTAP volumes with read-only access.
2. Set up a transition peer relationship between the 7-Mode system and the SVM.
3. Establishes a transition SnapMirror relationship (relationship of type TDP) between the 7-Mode volumes and ONTAP volumes.
4. Completes the baseline data copy transfer based on schedule inputs.
5. Performs scheduled incremental updates to the ONTAP volumes.

Apply configuration (precutover) phase

It is a best practice to run precutover operation a few days or weeks before the planned cutover window. This activity is to verify whether all the configurations are applied properly and whether any changes are required.

In this phase, configurations from the 7-Mode volumes are copied to ONTAP volumes.

There are two modes for the apply configuration (precutover) phase: **precutover read-only** and **precutover read/write**.

The precutover read/write mode is not supported when the project contains:

- SAN volumes and the target cluster is running Data ONTAP 8.3.1 or earlier

In this situation, the following configurations are not applied in the apply configuration (precutover) phase. Instead, they are applied during the cutover phase.

- SAN configurations
- Snapshot Schedule configurations
- SnapLock Compliance volumes

If the project contains SnapLock Compliance volumes, then the Snapshot Schedule configurations are not applied in the apply configuration (precutover) phase. Instead, these configurations are applied during the cutover phase.

Considerations for transitioning of SnapLock Compliance volumes

If the target cluster is running Data ONTAP 8.3.1 or earlier, and you want to run the apply configuration (precutover) operation in read/write mode for NAS volumes, then you must create separate projects for the NAS volumes and SAN volumes. This action is required because the precutover read/write mode is not supported if you have SAN volumes in your project.

If the project contains SnapLock Compliance volumes, and you want to run the apply configuration (precutover) operation in read/write mode for non-SnapLock Compliance volumes, then you must create separate projects for SnapLock Compliance volumes and non-SnapLock Compliance volumes. This action is required because the precutover read/write mode is not supported if you have SnapLock Compliance volumes in your project.

The tool performs the following steps in the **precutover read-only mode**:

1. Performs an incremental update from 7-Mode volumes to ONTAP volumes.
2. Breaks the SnapMirror relationship between 7-Mode volumes and ONTAP volumes.



For SnapLock Compliance volumes, the SnapMirror relationship between the 7-Mode volume and ONTAP volumes is not broken. The SnapMirror relationship is not broken because the SnapMirror resynchronization operation between 7-Mode and ONTAP volumes is not supported for SnapLock Compliance volumes.

3. Collects configurations from 7-Mode volumes, and applies the configurations to the ONTAP volumes and the SVM.
4. Configures the data LIFs on the SVM:
 - Existing 7-Mode IP addresses are created on the SVM in the administrative down state.
 - New IP addresses are created on the SVM in the administrative up state.
5. Resynchronizes the SnapMirror relationship between 7-Mode volumes and ONTAP volumes

The tool performs the following steps in the **precutover read/write mode**:

1. Performs an incremental update from 7-Mode volumes to ONTAP volumes.
2. Breaks the SnapMirror relationship between 7-Mode volumes and ONTAP volumes.
3. Collects configurations from 7-Mode volumes, and applying the configurations to the ONTAP volumes and the SVM.
4. Configures the data LIFs on the SVM:
 - Existing 7-Mode IP addresses are created on the SVM in the administrative down state.
 - New IP addresses are created on the SVM in the administrative up state.
5. Makes the ONTAP volumes available for read/write access.

After you apply the configuration, the ONTAP volumes are available for read/write access so that read/write data access can be tested on these volumes during apply configuration (precutover) testing. You can manually verify the configurations and data access in ONTAP.

6. Resynchronizes the ONTAP volumes when "finish testing" operation is triggered manually.

Storage cutover phase

The 7-Mode Transition Tool performs the following tasks in this phase:

1. Optional: Performs an on-demand SnapMirror update to reduce the downtime after cutover.
2. Manual: Disconnect client access from the 7-Mode system.
3. Performs a final SnapMirror update from 7-Mode volumes to ONTAP volumes.
4. Breaks and deletes the SnapMirror relationship between the 7-Mode volumes to ONTAP volumes, making the ONTAP volumes read/write.

If the selected volume is a SnapLock Compliance volume and the volume is the destination of a SnapMirror relationship, then the SnapMirror relationship between the 7-Mode volume and the ONTAP volume is deleted without a SnapMirror break operation. This action is performed to ensure that secondary ONTAP SnapLock Compliance volumes 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.

5. Applies Snapshot schedules configuration if:
 - The target cluster is running clustered Data ONTAP 8.3.0 or 8.3.1 and project contains SAN volumes.
 - The project contains SnapLock compliance volumes.
6. Applies SAN configurations, if the target cluster is running Data ONTAP 8.3.1 or earlier.
7. Applies quota configurations, if any.
8. Removes the existing 7-Mode IP addresses selected for transition from the 7-Mode system and brings the data LIFs on the SVM to the administrative up state.



SAN LIFs are not transitioned by the 7-Mode Transition Tool.

9. Optional: Takes the 7-Mode volumes offline.

Chain of Custody verification process for SnapLock volumes

You must perform the Chain of Custody verification operation. The tool performs the following operations when a Chain of Custody verification is initiated:

1. Enumerates all of the WORM files from 7-Mode volumes.
2. Calculates the fingerprint for each WORM file on the 7-Mode volumes (enumerated in the previous step) and calculates the fingerprint for the corresponding WORM file on the transitioned ONTAP volumes.
3. Generates a report with details about the number of files with matched and unmatched fingerprints, and the reason for the mismatch.



- The Chain of Custody verification operation is supported only for read-write SnapLock volumes that have file names with only ASCII characters.
- This operation can take a significant amount of time based on the number of files on the 7-Mode SnapLock volumes.

Post-transition steps

After the storage cutover phase finishes successfully and the transition is completed, you must perform some post-transition manual tasks:

1. Perform the required steps to configure features that were not transitioned or were partially transitioned, as listed in the precheck report.

For example, IPv6 and FPolicy must be configured manually after transition.

2. For SAN transition, reconfigure the hosts.

[SAN host transition and remediation](#)

3. Ensure that the SVM is ready to serve data to the clients by verifying the following:
 - The volumes on the SVM are online and read/write.
 - The IP addresses are up and reachable on the SVM.
4. Redirect client access to the ONTAP volumes.

Related information

[Migrating data and configuration from 7-Mode volumes](#)

How you transition volumes in a SnapMirror relationship

If you want to transition 7-Mode volumes that are in a SnapMirror relationship, the secondary volumes must be transitioned first. Then, a volume SnapMirror relationship is established between the 7-Mode primary volumes and ONTAP secondary volumes.

After transitioning the primary volumes, the 7-Mode Transition Tool establishes a volume SnapMirror relationship between ONTAP primary and secondary volumes.



The 7-Mode Transition Tool does not automatically transition SnapLock Compliance volumes that are in a SnapMirror relationship. All SnapLock Compliance volumes that are in a SnapMirror relationship must be transitioned as stand-alone volumes. After the primary and secondary SnapLock Compliance volumes are transitioned to ONTAP, you must manually perform the SnapMirror resynchronization operation between these volumes.

You can perform precheck, baseline copy, incremental transfers, and apply configuration (precutover) on the secondary and primary projects simultaneously; however, the storage cutover for the secondary project must be performed first.

Preparation phase

In this phase, the 7-Mode system, cluster, volumes, and IP addresses are selected. The 7-Mode Transition Tool performs the following tasks in this phase:

1. Adds 7-Mode storage system and volume information
2. Gathers information about 7-Mode source volumes and SnapMirror relationships:
 - For transitioning a secondary volume, collecting information about the 7-Mode primary system
 - For transitioning a primary volume, collecting information about the 7-Mode secondary system
3. Runs the transition precheck
4. Adds cluster, SVM, and aggregate information
5. Collects IP addresses that must be configured on the SVM:
 - Selecting IP addresses that exist on the 7-Mode system
 - Specifying new IP addresses that must be configured on the SVM



Transitioning iSCSI and FC LIFs (SAN) is not supported by the tool. You must manually configure the SAN LIFs on the SVM before transition.

6. Creates the data copy schedules for baseline and incremental transfers.
7. If the project contains SnapLock volumes, collects information about the read-write SnapLock volumes for which Chain of Custody verification is required and details about the ONTAP volume that stores the fingerprint data generated during the Chain of Custody verification operation.



The SnapLock Chain of Custody verification is supported only for read/write 7-Mode SnapLock volumes. It is not supported for read-only volumes. The SnapLock Chain of Custody verification is not supported for SnapLock volumes containing files that have names with non-ASCII characters.

8. Plans the configuration transition by selecting the 7-Mode configurations that must be transitioned to target SVM and target volumes.

You must not modify the objects (volumes, IP addresses, system information, and so on) on the controller after fixing errors and warnings that are reported by the precheck.

Data copy phase

In this phase, data from the 7-Mode volumes is copied to the ONTAP volumes. The 7-Mode Transition Tool performs the following tasks in this phase:

1. Creates the ONTAP volumes with read-only access
2. Set up a transition peer relationship between the 7-Mode system and the SVM
3. Establishes a SnapMirror relationship between the 7-Mode volumes and ONTAP volumes
4. Completes the baseline data transfer based on schedule inputs
5. Performs scheduled SnapMirror data copy updates to the ONTAP volumes

Apply configuration (precutover) phase

It is a best practice to run **Apply configuration** a few days or weeks before the planned cutover window. This precheck enables you to have enough time to verify that all of the configurations are applied properly and whether any changes are required.

In this phase, configurations from the 7-Mode volumes are copied to the ONTAP volumes.

There are two modes for the apply configuration (precutover) phase: precutover read-only and precutover read/write.

The precutover read/write mode is not supported when the project contains the following:

- SAN volumes and the target cluster is running Data ONTAP 8.3.1 or earlier

In this situation, the following configurations are not applied in the apply configuration (precutover) phase, instead they are applied during the cutover phase:

- SAN configurations
- Snapshot schedule configurations

- SnapLock Compliance volumes

If the project contains SnapLock Compliance volumes, then the Snapshot schedule configurations are not applied in the apply configuration (precutover) phase. Instead, these configurations are applied during the cutover phase.

Considerations for transitioning of SnapLock Compliance volumes.

If the target cluster is running Data ONTAP 8.3.1 or earlier and you want to run the apply configuration (precutover) operation in read/write mode for NAS volumes, then you must create separate projects for the NAS and SAN volumes. This action is required because the apply configuration (precutover) read/write mode is not supported if you have SAN volumes in your project.

If the project contains SnapLock Compliance volumes and you want to run the apply configuration (precutover) operation in read/write mode for non-SnapLock Compliance volumes, then you must create separate projects for SnapLock Compliance volumes and non-SnapLock Compliance volumes. This action is required because the apply configuration (precutover) read/write mode is not supported if you have SnapLock Compliance volumes in your project.

The following steps are performed by the tool in the **precutover read-only mode**:

1. Performs an incremental update from 7-Mode volumes to ONTAP volumes
2. Breaks the SnapMirror relationship between 7-Mode volumes and ONTAP volumes



For SnapLock Compliance volumes, the SnapMirror relationship between the 7-Mode volume and ONTAP volumes is not broken. This is because the SnapMirror resynchronization operation between 7-Mode and ONTAP volumes is not supported for SnapLock Compliance volumes.

3. Collects configurations from 7-Mode volumes and applying the configurations to the ONTAP volumes and SVM
4. Configures the data LIFs on the SVM:
 - Existing 7-Mode IP addresses are created on the SVM in the administrative down state.
 - New IP addresses are created on the SVM in the administrative up state.
5. Resynchronizes the SnapMirror relationship between 7-Mode volumes and ONTAP volumes

The following steps are performed in the **precutover read/write mode**:

1. Performs an incremental update from 7-Mode volumes to ONTAP volumes
2. Breaks the SnapMirror relationship between 7-Mode volumes and ONTAP volumes
3. Collects configurations from 7-Mode volumes and applying the configurations to the ONTAP volumes and SVM
4. Configures the data LIFs on the SVM:
 - Existing 7-Mode IP addresses are created on the SVM in the administrative down state.
 - New IP addresses are created on the SVM in the administrative up state.
5. Tests the read/write data access on the ONTAP volumes during apply configuration (precutover) testing

These ONTAP volumes will be available for read/write access after you apply the configuration. After you apply the configuration, the ONTAP volumes are available for read/write access so that read/write data


access can be tested on these volumes during apply configuration (precutover) testing.

6. Manual: Verifying the configurations and data access in ONTAP
7. Manual: Finish testing

The ONTAP volumes are resynchronized.

Storage cutover (secondary volumes) phase

The following illustration depicts the transition of a secondary volume:



| Phase | Steps |
|-------------------------------------|--|
| Storage cutover (secondary volumes) | <ol style="list-style-type: none"> 1. Transitioning the secondary volumes 2. Breaking and deleting SnapMirror relationship between the secondary volumes 3. Establishing a DR relationship between the 7-Mode primary and ONTAP secondary volumes |

The 7-Mode Transition Tool performs the following tasks in this phase:

1. Optional: Performs an on-demand SnapMirror update on the ONTAP secondary volumes
2. Manual: Disconnecting client access, if required
3. Performs a final SnapMirror update from the 7-Mode secondary volume to the ONTAP secondary volume
4. Breaks and deletes the SnapMirror relationship between the 7-Mode secondary volume and the ONTAP secondary volume, and making the destination volumes read/write
5. Applies the Snapshot schedules configuration, if the target cluster is running Data ONTAP 8.3.0 or 8.3.1 and the project contains SAN volumes
6. Applies SAN configurations, if the target cluster is running Data ONTAP 8.3.1 or earlier



All of the required igroups are created during this operation. For the secondary volumes, mapping LUNs to igroups is not supported during the cutover operation. You must manually map the secondary LUNs after completing the storage cutover operation of the primary volumes. However, for stand-alone volumes included in the secondary project, LUNs are mapped to the igroups during this operation.

7. Applies quota configurations, if any
8. Establishes a SnapMirror relationship between the volumes on the 7-Mode primary system and the ONTAP secondary volumes

The SnapMirror schedule that is used to update the SnapMirror relationships between the 7-Mode primary volumes and 7-Mode secondary volumes is applied to the SnapMirror relationships between the 7-Mode primary volumes and ONTAP secondary volumes.

9. Removes the existing 7-Mode IP addresses selected for transition from the 7-Mode system and bringing the data LIFs on the SVM to the administrative up state

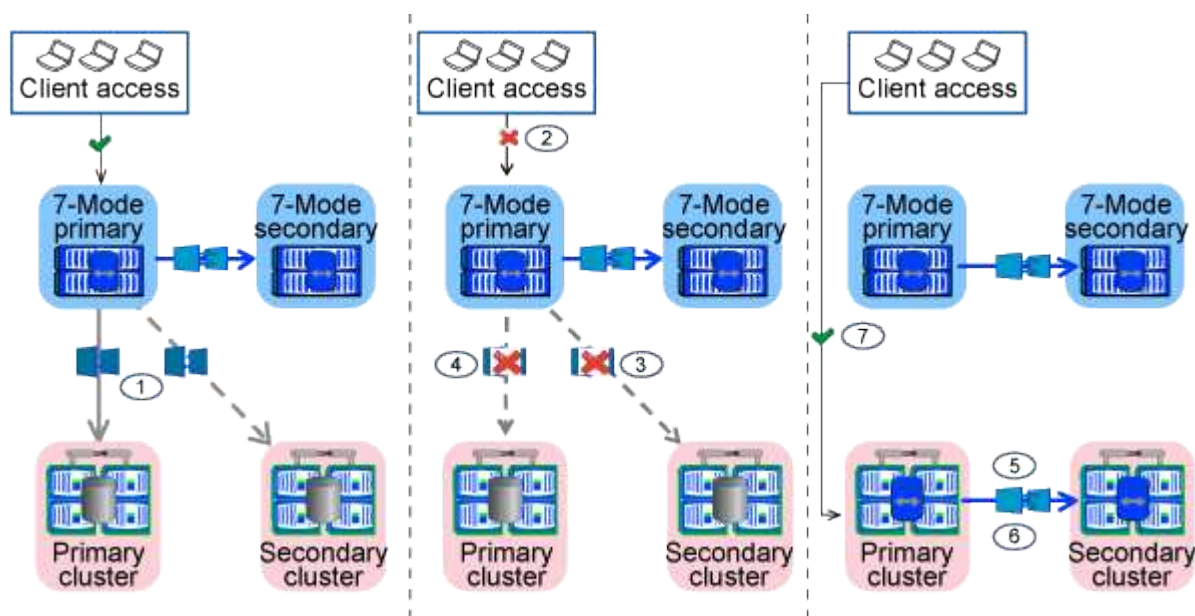


SAN LIFs are not transitioned by the 7-Mode Transition Tool.

10. Optional: Taking the 7-Mode volumes offline

Storage cutover (primary volumes) phase

The following illustration depicts the transition of a primary volume:



| Phase | Steps |
|-----------------------------------|--|
| Storage cutover (primary volumes) | <ol style="list-style-type: none"> 1. Transitioning the primary volumes 2. Disconnecting clients from the 7-Mode system (storage cutover) 3. Breaking and deleting the DR relationship between the 7-Mode primary and ONTAP secondary volumes 4. Breaking and deleting SnapMirror relationship between the primary volumes 5. Setting up an SVM peer relationship between the ONTAP primary and secondary volumes 6. Resynchronizing the SnapMirror relationship between ONTAP volumes 7. Enabling client access to ONTAP volumes |

The 7-Mode Transition Tool performs the following tasks in this phase:

1. Optional: Performs an on-demand SnapMirror update on the ONTAP secondary volumes
2. Manual: Disconnecting client access from the 7-Mode system
3. Performs a final incremental update from the 7-Mode primary volume and the ONTAP primary volume
4. Breaks and deletes the SnapMirror relationship between the 7-Mode primary volume and the ONTAP primary volume, and making the destination volumes read/write
5. Applies the Snapshot schedules configuration if the target cluster is running Data ONTAP 8.3.0 or 8.3.1 and the project contains SAN volumes
6. Applies SAN configurations, if the target cluster is running Data ONTAP 8.3.1 or earlier
7. Applies quota configurations, if any
8. Breaks and deletes the SnapMirror relationship between the 7-Mode primary volume and the ONTAP secondary volume
9. Setting up cluster peer and SVM peer relationships between the primary and secondary clusters
10. Setting up a SnapMirror relationship between the primary and secondary ONTAP volumes
11. Resynchronizes the SnapMirror relationship between the ONTAP volumes
12. Removes the existing 7-Mode IP addresses selected for transition from the 7-Mode system and bringing the data LIFs on the primary SVM to the administrative up state



SAN LIFs are not transitioned by the 7-Mode Transition Tool.

13. Optional: Taking the 7-Mode volumes offline

Chain of Custody verification process for SnapLock volumes

Perform the Chain of Custody verification operation.

1. Enumerates all of the WORM files from 7-Mode volumes

2. Calculates the fingerprint for each WORM file on the 7-Mode volumes (enumerated in the previous step) and calculates the fingerprint for the corresponding WORM file on the transitioned ONTAP volumes.
3. Generates a report with details about the number of files with matched and unmatched fingerprints, and the reason for the mismatch



- The Chain of Custody verification operation is supported only for read-write SnapLock volumes that have file names with only ASCII characters.
- This operation can take significant amount of time based on the number of files on the 7-Mode SnapLock volumes.

Post-transition steps

After the cutover phase is successfully and the transition is completed, you must perform the following post-transition tasks:

1. Perform any manual steps to transition features that were available on the 7-Mode system, but were not transitioned automatically to the SVM by the tool.
2. If the target cluster is running Data ONTAP 8.3.1 or earlier, you must map the secondary LUNs manually.
3. For SAN transitions, manually reconfigure the hosts.

[SAN host transition and remediation](#)

4. Ensure that the SVM is ready to serve data to the clients by verifying the following:
 - The volumes on the SVM are online and read/write.
 - The transitioned IP addresses are up and reachable on the SVM.
5. Redirect client access to the ONTAP volumes.

Related information

[Migrating data and configuration from 7-Mode volumes](#)

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 |
| 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 -sslsv3 -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.
```

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 | network route create <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 | network routing-groups route create <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.

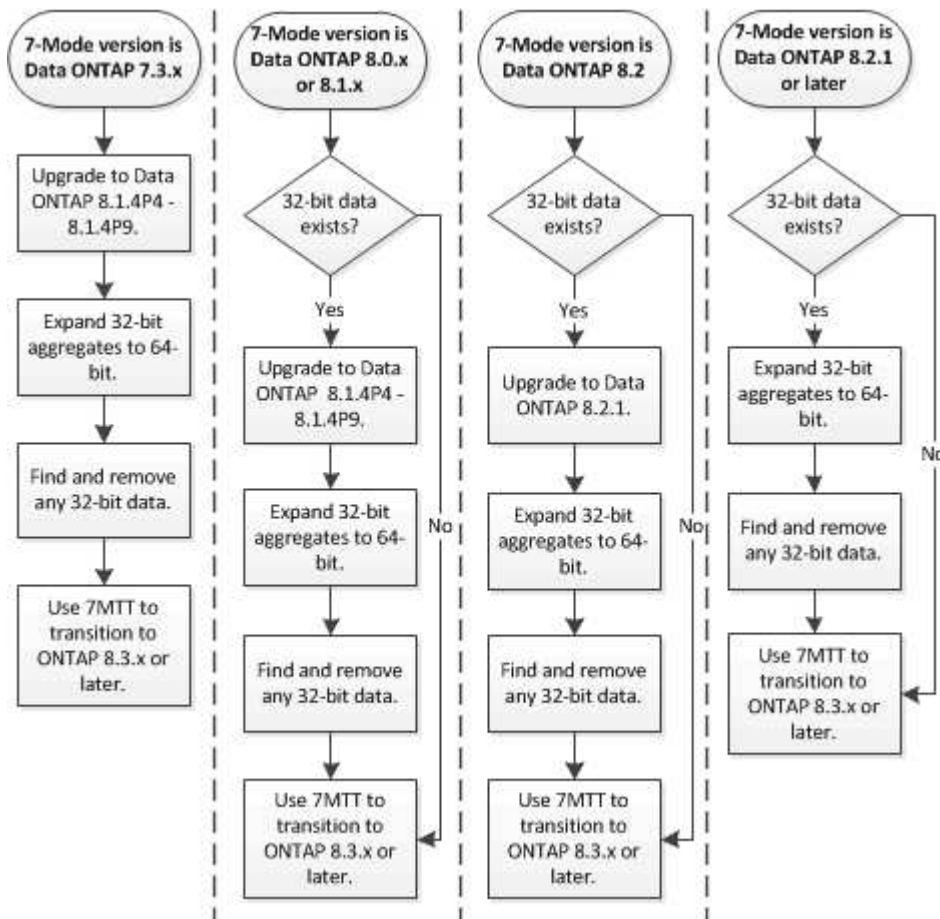
Preparing for transitioning to ONTAP 8.3 and later supported releases

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 <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
```

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`
 - `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`
- 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.

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 volsdt10 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/qtrees1
transition_volume_qtreemount:192.168.35.223:/vol/volstd10/qtrees1 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/qtree1 -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/qtree1  
[root@192.168.10.10]# ls transition_volume/qtree1  
[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 qtree1 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.
- 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 voll failed due to lack of space.

In this case, you must set the `filesys-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.

Migrating data and configuration from 7-Mode volumes


To migrate volumes or a volume SnapMirror relationship by using the 7-Mode Transition Tool, you must first configure projects, start a baseline copy, and complete the projects.

- The 7-Mode controllers and clusters that you want to include in the transition must be reachable from the Windows host where the tool is installed.
- You must have all administrator-level privileges for the controllers and clusters that you want to include in the transition project.
- The 7-Mode Transition Tool service must be running on the machine on which it is installed.

The service is set to automatic by default, and should start when you restart the machine.

- You should not perform assessment and migration operations on a controller simultaneously.
- You should not modify the objects (volumes, IP addresses, system information, and so on) on the 7-Mode controllers and clusters after fixing errors and warnings that are reported by precheck.
- You should avoid using multiple web interface sessions that are writing to the same SVM simultaneously to prevent undesired results.
- You should avoid modifying the controller and cluster passwords during the transition process.
- You should avoid using the **Back** and **Forward** browser buttons, as the tool does not support web browser navigation and might cause undesired results.
- You should avoid browser refresh while transition is in progress, because it might cause undesired results.

The following image illustrates the migration process:



Related information

[How you transition a stand-alone volume](#)

[How you transition volumes in a SnapMirror relationship](#)

Transition preparation checklist

Before you start transition, you should verify that you have met all of the prerequisites for transition.

ONTAP version requirements

| Item | Yes |
|--|-----|
| Supported 7-Mode version NetApp Interoperability Matrix Tool | |

| Item | Yes |
|---|-----|
| <p>Your cluster must be running one of the following Data ONTAP versions:</p> <ul style="list-style-type: none"> • Data ONTAP 8.2.x • Data ONTAP 8.3.x | |
| <p>You must transition to one of the following ONTAP releases:</p> <ul style="list-style-type: none"> • Using 7-Mode Transition Tool 3.3.3: <ul style="list-style-type: none"> ◦ ONTAP 9.8 or earlier supported releases • Using 7-Mode Transition Tool 3.3.2: <ul style="list-style-type: none"> ◦ ONTAP 9.7P2 or later 9.7 P release (earlier 9.7 releases are not supported) ◦ ONTAP 9.6P7 or later 9.6 P release (earlier 9.6 releases are not supported) ◦ ONTAP 9.5 or earlier ONTAP 9 release ◦ Clustered Data ONTAP 8.1.4P4 or later 8.x release • Using 7-Mode Transition Tool 3.3.1: <ul style="list-style-type: none"> ◦ ONTAP 9.5 or earlier ONTAP 9 release ◦ Clustered Data ONTAP 8.1.4P4 or later 8.x release | |

Licensing requirements

| Item | Yes |
|--|-----|
| SnapMirror license is enabled on the 7-Mode system | |
| SnapMirror licenses are enabled on the primary and secondary clusters for transitioning a volume SnapMirror relationship | |
| CIFS license is enabled on the cluster, if it is enabled on the 7-Mode system | |
| NFS license is enabled on the cluster, if it is enabled on the 7-Mode system | |
| iSCSI license is enabled on the cluster, if it is enabled on the 7-Mode system | |

| Item | Yes |
|---|-----|
| FC license is enabled on the cluster, if it is enabled on the 7-Mode system | |
| Other feature licenses, if available on the 7-Mode system, are added to the cluster | |

SnapMirror requirements on the 7-Mode system

| Item | Yes |
|--|-----|
| SnapMirror license | |
| <code>options snapmirror.enable on</code> | |
| <code>options interface.snapmirror.blocked ""</code> | |
| <p>Verify if one of the following is true:</p> <ul style="list-style-type: none"> • The <code>snapmirror.access</code> option is set to <code>all</code> • The <code>snapmirror.access</code> option is set to the IP addresses of all the intercluster LIFs on the cluster • If the <code>snapmirror.access</code> option is set to <code>legacy</code> and the <code>snapmirror.checkip.enable</code> option is <code>off</code>, the SVM name is added to the <code>/etc/snapmirror.allow</code> file • If the <code>snapmirror.access</code> option is set to <code>legacy</code> and the <code>snapmirror.checkip.enable</code> option is <code>on</code>, the IP addresses of the intercluster LIFs are added to the <code>/etc/snapmirror.allow</code> file | |

Volume settings on the 7-Mode system

| Item | Yes |
|--------------------------|-----|
| Volume is online | |
| Volume is not restricted | |

| Item | Yes |
|--|-----|
| <p>The following volume options are disabled:</p> <ul style="list-style-type: none"> • no_i2p • read_realloc • nvfail | |

Managing access to the cluster

| Item | Yes |
|---|-----|
| <p>SSL is enabled</p> <pre>system services web show</pre> | |
| <p>HTTPS is allowed on the cluster-management LIF</p> <pre>system services firewall policy show</pre> | |

Managing access to the 7-Mode system

| Item | Yes |
|---|-----|
| <p>HTTPS is enabled</p> <pre>options httpd.admin.ssl.enable on</pre> | |
| <p>SSL is enabled</p> <pre>secureadmin setup ssl</pre> <pre>options ssl.enable on</pre> | |
| <p>SSLv2 and SSLv3 are disabled</p> <pre>options ssl.v2.enable off</pre> <pre>options ssl.v3.enable off</pre> | |

Networking requirements

| Item | Yes |
|--|-----|
| <p>Cluster is reachable using the cluster-management LIF</p> | |

| Item | Yes |
|--|-----|
| One or more intercluster LIFs are set up on each node of the cluster For multipathing, two intercluster LIFs are required on each node | |
| Static routes are created for the intercluster LIFs | |
| 7-Mode system and cluster are reachable from the Windows system on which 7-Mode Transition Tool is installed | |
| NTP server is configured and the 7-Mode system time is synchronized with the cluster time | |

Port requirements

| Item | Yes |
|--|-----|
| 7-Mode system <ul style="list-style-type: none"> • 10565/TCP • 10566/TCP • 10567/TCP • 10568/TCP • 10569/TCP • 10670/TCP • 80/TCP • 443/TCP | |
| Cluster <ul style="list-style-type: none"> • 10565/TCP • 10566/TCP • 10567/TCP • 10568/TCP • 10569/TCP • 10670/TCP • 11105/TCP • 80/TCP • 443/TCP | |

NFS requirements

| Item | Yes |
|--|-----|
| NFS license is added to the cluster | |
| DNS entry must be configured for AD domain on the SVM | |
| NFS is added to the list of allowed protocols for the SVM | |
| Clock skews between KDC and the cluster is less than or equal to 5 minutes | |

CIFS requirements

| Item | Yes |
|---|-----|
| CIFS license is added to the cluster | |
| If MultiStore license is enabled, CIFS must be added to the list of allowed protocols for the vFiler unit that owns the transitioning volumes | |
| CIFS is set up and running on the 7-Mode system | |
| Authentication type in 7-Mode for CIFS is Active Directory (AD) or Workgroup | |
| CIFS is added to the list of allowed protocols for the SVM | |
| DNS is configured for the SVM | |
| CIFS server is configured for the SVM | |
| CIFS is running on the SVM | |

Related information

[Preparing for copy-based transition](#)

Adding controllers and clusters

Before you start the transition, you must add the 7-Mode controllers and clusters that are required for the transition. The 7-Mode controllers that are included for assessment are automatically added for migration.

- The 7-Mode controller and cluster information that you provide is not persistent.

If the 7-Mode Transition Tool service is restarted, the tool prompts you for information about controllers and clusters that are part of active projects. You must provide the same host name that you provided for your system when you created the project.

- If a 7-Mode controller is part of an HA pair, the tool does not request for credentials of the HA partner of the 7-Mode controller (unless the HA partner is part of another active project.)


1. From the top pane, click **Storage Systems**.
2. In the **Hostname** field, enter the FQDN or IP address of the 7-Mode controller or the ONTAP system.

For a cluster, you can specify the IP address or FQDN of the cluster-management interface. For a 7-Mode controller, you must specify the IP address of the default vFiler unit, because the IP addresses of individual vFiler units are not accepted.

Steps

1. Enter the administrator credentials for the specified host, and then click **Add**.

The 7-Mode controllers are added to the “7-Mode Controllers” table and clusters are added to the “Clustered Data ONTAP Systems” table.

2. Repeat Steps 2 and 3 to add all of the controllers and clusters that you require for the transition.
3. If the Status column indicates that the credentials of the system are missing or the credentials have changed from what was initially entered in the tool, click the  icon, and then enter the credentials again.

Creating a transition project

Creating a transition project includes selecting and mapping 7-Mode volumes to the storage virtual machine (SVM), mapping interfaces, and creating data copy schedules for SnapMirror relationships.

You must have created the required SVM on the cluster.

All of the volumes within a project are migrated to the same SVM. If you want to migrate the volumes to different SVMs, you must create multiple projects.

If the target cluster is running Data ONTAP 8.3.1 or earlier and you want to run the precutover operation in read/write mode for NAS volumes, then you must create separate projects for the NAS volumes and SAN volumes. This action is required because the precutover read/write mode is not supported if you have SAN volumes in your project.

If the project contains SnapLock Compliance volumes and you want to run the precutover operation in read/write mode for non-SnapLock Compliance volumes, then you must create separate projects for SnapLock Compliance volumes and non-SnapLock Compliance volumes. This action is required because the precutover read/write mode is not supported if you have SnapLock Compliance volumes in your project.

Steps

1. Select the **Copy-Based Transition** migration method from the homepage, and then click **Start Planning**.

If the controller and cluster that are required by the new project have not been added, you can enter the details in the Enter Device Credentials pane.

2. Verify that all of the required Data ONTAP operating in 7-Mode systems and ONTAP systems are added to the tool, and then click **Next**.

The Select Source Volume page appears.

3. Select the 7-Mode volumes that you want to transition.
 - a. From the 7-Mode Controller pane, select the 7-Mode controller or the vFiler unit from which you want to add volumes.
 - b. Add the volumes that you want to include in the project group:

| If you want to transition... | Then... |
|--------------------------------|--|
| Stand-alone volumes | <p>Select Transition as stand-alone for the volumes that you want to transition.</p> <p>A stand-alone project is created if you select the first volume from this column.</p> |
| Volume SnapMirror relationship | <p>i. Select Transition with SnapMirror Relationship for all of the primary volumes.</p> <p>Two projects are created: a primary project for the primary volumes and a secondary project.</p> <p>ii. Optional: If the secondary controller is not included in the project, enter the details for the controller in the Need additional storage system credentials dialog box.</p> |

If you have at least one LUN in your volume, the volume type is shown as SAN.

The hyperlink that is provided on the volume name opens a dialog box that lists the qtrees and LUNs in the volume and their attributes.



It is a best practice to have all of the volumes within a single project to be of the same definition (stand-alone, primary, or secondary). For example, a project should contain all stand-alone volumes rather than a mix of stand-alone and SnapMirror relationships.

- c. After you have selected all of the volumes that you want to include in the project, click **Create Project and Continue**, enter the project name and project group details from the dialog box that appears, and then click **Save** to create the project.
4. Select the 7-Mode IP address and multipath IP address to be used for SnapMirror data copy.
 - a. Enter the 7-Mode data copy IP address.

By default, this field is prepopulated with the management IP address of the 7-Mode system. If required, you can change this IP address to any valid IPv4 address with data copy permission.
 - b. If you want to use multiple paths for load balancing the data transfers, enter an IP address in the IP Configuration pane, and then click **Next**.
5. From the Select SVM page, select the target cluster and SVM and follow these steps:

- a. Select the target cluster by clicking on the cluster name in the Select a Clustered Data ONTAP System drop-down list.

The SVMs are loaded in the Select SVM pane.

- b. Select the target SVM to transition the volumes from the Select SVM pane.
- c. Click **Next**.

For transitioning 7-Mode volumes to a MetroCluster configuration in ONTAP, the SVM subtype must be `sync-source`.

+ If you select an SVM that belongs to clustered Data ONTAP 8.2, a dialog box is displayed to confirm whether local users and groups or CIFS shares or files are configured on the 7-Mode storage system. The 7-Mode Transition Tool does not support the transition of local users and groups to clustered Data ONTAP 8.2. If you have local users and groups, you can select an SVM that belongs to ONTAP 8.2.1 and later supported releases.

6. In the SVM audit logs destination path dialog box, enter a path on the destination SVM to enable transition of the audit configuration from the 7-Mode storage system.

This path is used to save the audit logs in the ONTAP system.

7. From the Map Volumes page, select the target volumes for transition to map each source volume to the required aggregate.
 - a. From the Map Origin Volumes to Aggregates on Target Cluster pane, select the aggregates to which the 7-Mode volumes must be copied.
 - b. To change the name of the target volume on the cluster, enter a different name in the **Target Volume** field.
 - c. Click **Next**.

If all of the volumes and qtrees that are included in the project are configured to serve only NFS requests, then you do not have to provide the audit path because the audit configuration is not transitioned (even if you provide the audit path, this input is ignored) .

8. From the Network Configuration pane, provide information about the LIFs that must be created on the SVM.



FC and iSCSI LIFs cannot be transitioned. You must manually create them on the SVM.

| If you want to... | Then... |
|--|---|
| Transition an existing 7-Mode IP address | <ol style="list-style-type: none"> a. Click Select 7-Mode LIF. b. Select the required 7-Mode IP addresses, and provide target node and target port details. c. Click Save. |
| Create a new LIF | <ol style="list-style-type: none"> a. Click Add New LIF. b. In the dialog box that appears, enter the details for the new LIF. c. Click Save. |

To provide network connectivity after a successful transition, you must 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.

9. After you add all the required IP addresses, click **Next**.
10. In the Configure Schedule page, configure the data copy schedules for baseline and incremental transfers, the number of concurrent volume SnapMirror transfers, and the throttle limit for the SnapMirror transfers for transition.

You can provide data copy schedules and a throttle limit to effectively manage your DR and transition data copy operations. You can create multiple schedules, with a maximum of seven schedules for each project. For example, you can create customized schedules for weekdays and weekends.



The schedules are effective based on the source 7-Mode controller time zone.

- a. In the Configure Schedule pane, click **Create Schedule**.
- b. In the Create Data Copy Schedule dialog box, enter a name for the new schedule.
- c. In the Recurring Days pane, select **Daily** or **Select Days** to specify the days on which the data copy operations should run.
- d. In the Time Interval pane, specify the **Start Time** and **Duration** for the data transfers.
- e. In the Time Interval pane, either specify the **Update Frequency** for the incremental transfers or select **Continuous Update**.

If you enable continuous updates, the updates start with a minimum delay of 5 minutes, depending on the availability of concurrent SnapMirror transfers.

- f. In the Parameters for Transition Data Copy Operations (based on Volume SnapMirror) pane, specify the maximum number of concurrent volume SnapMirror transfers (as a percentage of available SnapMirror transfers at run time and as a number) and the throttle limit (maximum bandwidth for all of the volumes in the project).



The default values that are provided in the fields are the recommended values. When changing the default values, you must analyze the 7-Mode SnapMirror schedules and ensure that the values that you provide do not affect these schedules.

- g. Click **Create**.

The new schedule is added to the Transition Schedule pane.

- h. After you add all of the required data copy schedules, click **Next**.
11. If you have SnapLock volumes to transition, plan the volumes that require Chain of Custody verification after transition.
 - a. Select the source SnapLock volumes that require Chain of Custody verification.

The Chain of Custody verification process is supported only for read/write 7-Mode SnapLock volumes and is not supported for read-only volumes. Only SnapLock volumes that have file names with ASCII characters are supported for Chain of Custody verification.

- b. Provide details about the ONTAP volume that will be used to store the fingerprint data generated during the Chain of Custody verification operation.

The ONTAP volume must already exist on the specified SVM.

- c. Click **Next**.

Related information

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

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

[Managing SnapMirror transfers and schedule](#)

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

[Managing logical interfaces](#)

[Removing volumes from a project](#)

Customizing the transition of 7-Mode configurations

When planning the transition of configurations from 7-Mode to ONTAP, you can customize the configuration transition in two ways. You can ignore or skip the transition of one or more configurations. You can consolidate the 7-Mode NFS export rules, and then reuse an existing NFS export policy and Snapshot policy on the target SVM.

You must perform this task before you apply the configuration (precutover) phase. This is because after this phase, the Plan Configuration pane is disabled for any modification. You use the command-line interface (CLI) of the 7-Mode Transition Tool for excluding the configurations that are applied during the cutover phase.

The 7-Mode Transition Tool does not perform prechecks for the configuration that is excluded.

By default, all 7-Mode configurations are selected for transition.

It is a best practice to run the prechecks with all configurations first, and then exclude one or more configurations in the subsequent run of the prechecks. This helps you to understand which configurations are excluded from transition and which prechecks are skipped subsequently.

Steps

1. From the Plan Configuration page, select the following options from the **SVM Configuration** pane:
 - For excluding the transition of configurations, clear the check box for those configurations.
 - For consolidating similar 7-Mode NFS export rules to a single export policy in ONTAP, which can then be applied to the transitioned volume or qtree, select the **Consolidate NFS Export Policies on 7-Mode** check box.
 - For reusing an existing NFS export policy on the SVM that matches the export policy that will be created by the tool, which can then be applied to the transitioned volumes or qtrees, select the **Reuse Export Policies of SVM** check box.
 - For consolidating similar 7-Mode Snapshot schedules to a single Snapshot policy in ONTAP, which can then be applied to the transitioned volume, select the **Consolidate 7-Mode Snapshot Policies** check box.

- For reusing an existing Snapshot policy on the SVM that matches the Snapshot policy that will be created by the tool, which can then be applied to the transitioned volumes, select the **Reuse Snapshot Policies of SVM** check box.

2. Click **Save and go to Dashboard**.

Related information

[Supported and unsupported CIFS configurations for transition to ONTAP](#)

[NFS transition: supported and unsupported configurations, and required manual steps](#)

[Name services transition: supported and unsupported configurations, and required manual steps](#)

[SAN transition: supported and unsupported configurations, and required manual steps](#)

[Examples of consolidating NFS export rules and Snapshot schedules for transition](#)

[Configurations that can be excluded from transition](#)

Running prechecks

You can run prechecks to identify any issues before you start a transition. Prechecks verify that the 7-Mode sources, ONTAP targets, and configurations are valid for your transition. You can run prechecks any number of times.

The prechecks run more than 200 different checks. For example, the tool checks for items such as if volumes are online and network access exists between the systems.

Steps

1. From Dashboard, select the project for which you want to run the prechecks.
2. Click **Run Prechecks**.

After the prechecks are complete, the result summary is displayed in the dialog box.



The prechecks usually take only a few minutes to run, but the duration of the precheck phase depends on the number and type of errors or warnings that you resolve.

3. Choose an option under **Apply Type Filter** to filter the results:
 - To view all messages related to security, select **Error**, **Warning**, **Informational**, and **Security Only**.
 - To view all error messages related to security, select **Error** and **Security Only**.
 - To view all warning messages related to security, select **Warning** and **Security Only**.
 - To view all informational messages related to security, select **Informational** and **Security Only**.
4. To save the raw results in comma-separated values (CSV) format and export the results, click **Save As CSV**.

You can view the transition operations that have been performed during the transition along with the operation type, status, start time, end time, and results in the Operation History tab on the Dashboard pane.

You must resolve all the errors detected by the prechecks before you start data copy. It is also a best practice

to resolve all warnings prior to proceeding with the migration process. Resolution can be resolving the source issue of the warning message, implementing a workaround, or accepting the result of the issue.

Severity levels for precheck messages

You can verify whether the 7-Mode volumes can be transitioned by running the transition precheck operation. Transition precheck reports all the transition issues. Transition issues are assigned different severity levels, depending on the impact of the issue on the transition process.

The issues detected by the prechecks are classified into the following categories:

- **Error**

Configurations that cannot be transitioned.

You cannot continue the transition if there is even one error. The following are a few example configurations on the 7-Mode system that cause an error:

- Traditional volumes
- SnapLock volumes
- Offline volumes

- **Warning**

Configurations that can cause minor problems after transition.

Features that are supported in ONTAP, but are not transitioned by the 7-Mode Transition Tool, also generate a warning message. You can continue the transition with these warnings. However, after the transition you might lose some of these configurations or might have to complete some manual tasks for enabling these configurations in ONTAP.

The following are a few example configurations on the 7-Mode system that generate a warning:

- IPv6
- NFSv2
- NDMP configurations
- Interface groups and VLANs
- Routing Information Protocol (RIP)

- **Information**

Configurations that have been successfully transitioned.

Starting baseline data copy

After you create a project and complete the precheck operation, you must initiate data copy from the 7-Mode volumes to ONTAP. You can start baseline data copy operation for individual projects. You should stop unnecessary system processes and network activity during the data copy.

You must have created at least one data copy schedule.

You can estimate the time to complete baseline transfers and evaluate the performance achieved by volume SnapMirror transfers in your environment by performing a test migration. The following are some of the factors that can affect performance:

- Transition data copy schedule options selected

This schedule controls both the maximum number of SnapMirror concurrent transfers and the maximum bandwidth to be used for the transfers.

- Maximum number of concurrent volume SnapMirror transfers supported by the 7-Mode source controllers
- Network bandwidth between the 7-Mode source and ONTAP destination controllers

Network traffic that is unrelated to the migration activity must be minimized so that the throughput is maximized and response time is minimized between the source and destination systems.

- Performance capabilities of both the source and destination controllers

The source and destination systems should have optimum CPU utilization and memory available.

- Number of 7-Mode volume SnapMirror transfers occurring during the data copy

Steps

1. From Dashboard, select the project for which you want to start the baseline data copy.
2. Click **Start Baseline**.

The precheck is run once again in the background, and if no errors are detected, the baseline transfer is started based on the data copy schedule. The Operation Progress dialog box displays the information about the status of the precheck operations run during the baseline data copy.

3. Click the **Volumes** tab to view the status and progress of the baseline transfer.

To view the detailed SnapMirror details of each volume, you can click **View Transition Details**. The number of concurrent SnapMirror transfers is based on the input provided in the schedule that is currently active. You can track the active schedule from the Data Copy Schedule tab on Dashboard.

After the baseline data copy operation is completed, the incremental SnapMirror updates start based on the schedule provided while creating the project.

Related information

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

Applying 7-Mode configurations

After the baseline data copy is completed, you can copy and apply all configurations from the 7-Mode system (including protocols and services configuration) to the ONTAP volumes. If the target cluster is running any version from ONTAP 8.3.2 and later supported releases, SAN configuration is transitioned in this phase.

If you are transitioning SAN volumes, you must have created at least one data LIF of the appropriate protocol (iSCSI or FC) for every node in the cluster.

- The configurations are applied in the apply configuration (precutover) phase, which has two modes: precutover read-only mode and precutover read/write mode.

The precutover read/write mode is not supported when the project contains:

- SAN volumes and the target cluster is running Data ONTAP 8.3.1 or earlier. In this situation, the following configurations are not applied in the precutover phase, instead they are applied during the cutover phase:
 - SAN configurations
 - Snapshot Schedule configurations
- SnapLock Compliance volumes.

If the project contains SnapLock Compliance volumes, then the Snapshot Schedule configurations are not applied in the precutover phase, instead these configurations are applied during the cutover phase.

See [Considerations for transitioning of SnapLock Compliance volumes](#).

Steps

1. From the Dashboard, select the project.
2. Apply the configurations:

| If you want to apply all configurations in... | Then... |
|---|---|
| Read-only mode | Click Apply Configuration . |
| Read/write mode | <ol style="list-style-type: none"> a. Select the Test Mode check box. b. Click Apply Configuration. The ONTAP volumes are made read/write and you can test the configurations and data access operations. c. Select Apply configuration in test mode in the Apply Configuration (Precutover) dialog box. |

3. Select the **Customize the number of concurrent SnapMirror transfers and Throttle limit for this operation** check box to specify the number of SnapMirror data copy operations and throttle limit:
 - a. Enter the maximum number of concurrent SnapMirror transfers to run during transition.
 - b. Enter the percentage of available streams that can be used for SnapMirror transfers.

By default, the tool uses 50% of the available volume SnapMirror transfers.
 - c. Either enter a throttle limit or select **Maximum** to use the maximum bandwidth.

By default, the tool uses maximum throttle for configuration transition.
4. Select the **Transition Kerberos Configuration** check box to provide UNIX-based or Microsoft AD based Kerberos server configuration details for transition.



This option is enabled only when Kerberos is configured on the source 7-Mode storage system.

- a. Enter the Kerberos server details such as the host name, IP address, user name, and password.



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

5. Click **Continue**.

The Operation Progress dialog box is displayed, and the copy configuration operation is started.

6. If the configuration transition is performed in read/write mode, click **Finish Testing** after the testing and verification of the configurations is complete.

This mode should be used only for testing purposes. All data written to the cluster on the volumes being migrated during test mode is lost.

The tool reestablishes the SnapMirror relationship and resynchronizes (based on the active schedule for that project at that time) the ONTAP volumes. Any data written to the 7-Mode is resynchronized with the ONTAP volumes.



For a successful resynchronization, a common Snapshot copy must exist between the 7-Mode and clustered Data ONTAP volumes. You should not manually delete the common Snapshot copy; otherwise, resynchronization fails.

The 7-Mode IP addresses remain operational. The LIFs are configured on the storage virtual machine (SVM) in the following ways:

- Existing 7-Mode IP addresses are created in the administrative down state.

During the storage cutover, these IP addresses are removed from the 7-Mode system and the corresponding storage virtual machine (SVM) LIFs are brought to the administrative up state. If you select the precutover read/write mode, you must use a different LIF to gain access to the volumes being migrated to the cluster.

- New IP addresses are created in the administrative up state.

If you select the precutover read/write mode, these LIFs can be used for testing access to the volumes being migrated in the cluster.

Related information

[Managing logical interfaces](#)

[Considerations for quotas](#)

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

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 pwwn 21:00:00:c0:dd:19:4b:f3
member pwwn 20:07:00:a0:98:32:99:07
member pwwn 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.

Performing on-demand SnapMirror updates

You can perform SnapMirror incremental updates for all the volumes before the cutover operation to reduce the cutover time.

- You cannot perform on-demand SnapMirror updates when incremental data transfers are scheduled after baseline data copy and after precutover operation.
- This is an optional task.

1. Click **Update Now** to perform a manual SnapMirror update.

The Transition Update dialog box is displayed, where you can choose to customize the number of SnapMirror transfers and throttle limit for this operation.

2. Select the **Customize the number of concurrent SnapMirror transfers and Throttle limit for this**

operation check box to specify the number of SnapMirror data copy operations and throttle limit.

- a. Enter the maximum number of concurrent SnapMirror transfers to run during transition.
- b. Enter the percentage of available streams that the tool can use for SnapMirror transfers.

By default, the tool uses 50% of the available volume SnapMirror transfers.

- c. Enter the throttle limit to use the maximum bandwidth.

By default, the tool uses maximum throttle for configuration transition.

3. Click **Continue**.

Related information

[Starting baseline data copy](#)

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

Completing a transition project

You can complete a transition by completing the individual projects. Because this operation is disruptive, you should evaluate when to run it. When transitioning volumes in a SnapMirror relationship, the secondary project must be completed before completing the transition of the primary project.

The storage cutover is completed in a few minutes. The time required for the clients to remount the data varies. The timing of the storage cutover or outage window depends on the following factors:

- Final update

The final update of the data depends on the amount of change in the source data since the last update. Incremental transfers minimize the amount of data that has to be transferred during cutover.

- Reconnecting clients

If updates are required for each client to connect to the cluster, the number of clients that have to be updated determines the cutover time.

Outages apply only to the volumes that are being migrated. You do not need to shut down the entire source 7-Mode storage system. Volumes on the source system that are not being migrated can remain online and accessible.

1. From the Migration Dashboard, select the project that you want to complete.
2. Disconnect client access manually.
3. Click **Complete Transition**.

- a. If you want to keep the 7-Mode source volumes online after the transition, clear the **Take source volumes offline after transition** checkbox.

By default, this option is selected, and the source volumes are taken offline.

- b. If you have selected SnapLock volumes for Chain of Custody verification, select the **I understand that**

I must not take 7-Mode SnapLock volumes offline during Chain of Custody verification checkbox to keep the SnapLock volumes online after transition.

- c. If you have selected the transition of a SnapMirror relationship between clusters that are running ONTAP 9.3 or later supported releases, select the **I understand that I must manually convert SnapMirror relationship type from data_protection to extended_data_protection** checkbox.
- d. Select the **Customize the number of concurrent SnapMirror transfers and Throttle limit for this operation** checkbox to specify the number of SnapMirror data copy operations and the throttle limit.
- e. Click **Continue**.

The results of the cutover operation are displayed.

The 7-Mode IP addresses selected for the transition are unconfigured from the 7-Mode storage system, and the associated LIFs created before the cutover are brought to the administrative up state. The 7-Mode volumes are offline.

From the cluster, run the `vserver check lif-multitenancy run` command to verify that the name servers are reachable by using the transitioned LIFs.



If you have created a new LIF, the users and applications of the transitioned volumes must be remapped to the drives by using the new IP addresses and ports after all of the projects have been completed.

If you have completed the transition of a SnapMirror relationship between clusters that are running ONTAP 9.3 or later supported releases, you must convert the SnapMirror relationship from type DP to type XDP.

Data protection

Related information

[Guidelines for deciding when to perform cutover](#)

Completing the Chain of Custody verification

If one or more SnapLock volumes are selected for Chain of Custody verification, then you must perform the Chain of Custody operation to generate a Chain of Custody report.

You must have completed the transition of the project.

SnapLock Chain of Custody operation is supported for volumes with files that have file names with only ASCII characters.

1. From the Migration Dashboard, click **Start Chain of Custody**.

If you want to keep the 7-Mode SnapLock volumes online after the Chain of Custody verification, you should clear the **Take 7-Mode SnapLock volumes selected for Chain of Custody verification offline after Chain of Custody verification operation is completed** check box.

2. Click **Continue**.

The Chain of Custody verification operation is initiated. This operation can take a significant amount of time based on the number of files on the SnapLock volumes. You can click **Run in Background** to perform the operation in the background.

You can track the progress of the Chain of Custody verification operation by clicking the SnapLock Chain of Custody tab in the Migration Dashboard window. This tab displays per volume progress of the Chain of Custody operation.

3. After the Chain of Custody operation is complete, click **Download Report** in the SnapLock Chain of Custody tab to download the Chain of Custody verification report.

The Chain of Custody verification report contains details about whether the SnapLock Chain of Custody verification succeeded. The report shows the total file count and the number of non-WORM files in each of the 7-Mode SnapLock volumes that are selected for the Chain of Custody operation. You can also verify the number of files for which the fingerprints matched and unmatched. The report also shows the number of WORM files for which the Chain of Custody verification failed and reason for the failure.

Transitioning volumes by using the 7-Mode Transition Tool CLI

The 7-Mode Transition Tool provides commands that enable you to transition 7-Mode volumes. If you have the 7-Mode Transition Tool installed on a Linux system, you must use the CLI to perform the transition.

The tool collects the 7-Mode storage system's volume and IP address information, verifies the 7-Mode configuration, and interacts with the cluster to apply the configuration on the SVM.


Scenarios for selecting volumes in a project

Planning a project composition and deciding which 7-Mode volumes to transition in a single transition operation can help in reducing the downtime for applications. Understanding some sample scenarios can help you while creating projects in your environment.

Bidirectional SnapMirror relationship

To transition a bidirectional SnapMirror relationship, you must create different transition projects for each volume.

For example, as shown in the illustration below, consider a volume SnapMirror relationship that exists between vf1:vol1 (primary) on system1 and vf2:vol1 (secondary) on system2. Similarly, another volume SnapMirror relationship exists between vf2:vol2 (primary) on system2 and vf1:vol2 (secondary) on system1.



You cannot combine the volumes vf1:vol1 and vf1:vol2 in one transition project. Similarly, you cannot combine the volumes vf2:vol1 and vf2:vol2 in one transition project. You must create separate transition projects for each volume to transition the volume SnapMirror relationship.

Volumes with CIFS configuration

You must group volumes that have CIFS configuration in one project so that all the CIFS-related configuration is transitioned completely to the SVM.

For example, if 10 volumes in a 7-Mode system or a vFiler unit have associated CIFS shares, home-directory search path, and audit configuration, these 10 volumes must be transitioned in one project. This ensures that all the volumes and CIFS configuration are completely applied on the SVM after transition.

SnapMirror relationship between a primary volume and multiple secondary volumes

If a SnapMirror relationship exists between a primary volume and multiple secondary volumes and if all the secondary volumes are on the same 7-Mode controller, you can create a secondary project to group all the secondary volumes and complete the transition of all secondary volumes in that project. You can then create a primary project to transition the primary volume and complete the SnapMirror transition.

Transitioning volumes by using the 7-Mode Transition Tool CLI

The 7-Mode Transition Tool provides commands that enable you to transition 7-Mode volumes. If you have the 7-Mode Transition Tool installed on a Linux system, you must use the CLI to perform the transition.

The tool collects the 7-Mode storage system's volume and IP address information, verifies the 7-Mode configuration, and interacts with the cluster to apply the configuration on the SVM.

Creating a transition project

You can create a transition project that identifies the 7-Mode objects to be transitioned and how they map to the ONTAP objects. While creating the project, the 7-Mode Transition Tool collects and saves the information about the transition objects. You can

use this project information to perform the transition in stages.

You must have gathered the following information:

- Project type: Possible values are stand-alone, primary, or secondary
- FQDN or IP address of the 7-Mode storage system or the default vFiler unit, if MultiStore is licensed
- Administrative user name and password of the 7-Mode storage system
- List of 7-Mode volumes
- FQDN of the cluster or IP address of the cluster-management LIF
- SVM name

Steps

1. From the **Start** menu, click **All Programs > NetApp 7-Mode Transition Tool > NetApp 7-Mode Transition Tool (CLI) 3.0** to open the 7-Mode Transition Tool command prompt.
2. Add the user name and password of the systems that are a part of the transition project:

```
transition credentials add -h host_name -u user_name
```

host_name is FQDN or IP address of the 7-Mode system or cluster.

user_name is the administrative user name for the system.



You must repeat this step for each storage system that you want to add to the transition projects.

```
7-Mode Transition Tool>transition credentials add -h system1.example.com  
-u root  
Enter password for 'root@system1.example.com':
```

3. Create a transition project:

```
transition cbt create -p project_name -t project_type -n 7-Mode_system -c  
data-copy-ipaddress [-f vfiler_name] [-h cluster -v vserver_name
```

project_name is the name of the transition project.

project_type is the project type. Valid values are standalone, secondary, or primary.

7-Mode_system is the FQDN or IP address of the 7-Mode system.

cluster is the FQDN of the cluster or IP address of the cluster-management LIF.



The host names of your 7-Mode system and cluster must be the same as those you had provided while creating the project.

data-copy-ipaddress is the IP address over which data is copied.

vfiler_name is the name of the nondefault vFiler unit.

vserver_name is the name of the SVM.



For a new SVM, you must provide the aggregate name of the root volume of the SVM with the `-g` option.

```
7-Mode Transition Tool>transition cbt create -p sample_project -t
standalone -n system1.example.com -n 10.238.55.33 -h
cluster1.example.com -v vs2
```

The transition project moves to the preparation state.

4. Add the 7-Mode volume and ONTAP volume to the project:

```
transition cbt volumepair add -p project_name -v 7-Mode_volume -c
vserver_volume -g aggr_name -t
```

7-Mode_volume is the 7-Mode volume name.

vserver_volume is the ONTAP volume on the SVM.



You can also add an ONTAP volume for which the baseline transfer has been completed. When you add such a volume in the preparation phase, only incremental transfers are performed during the data copy phase.

aggr_name is the aggregate on which the ONTAP volume *vserver_volume* is created during the start operation.

`-t` is used when a volume which is in SnapMirror relationships must be transitioned as standalone volume.



When a volume is transitioned as standalone volume, SnapMirror relationships are not preserved during the transition.

For SnapLock Compliance volumes that are in SnapMirror relationship, the transition of SnapMirror relationships for SnapLock Compliance volumes is not automated by the 7-Mode Transition Tool. All SnapLock Compliance volumes that are in SnapMirror relationship must be transitioned as standalone volumes (by specifying the flag `-t`) by including the volumes in standalone, primary, or secondary projects. You must then manually perform the SnapMirror resynchronization operation between these volumes after the transition.

See [Considerations for transitioning of SnapLock Compliance volumes](#)

```
7-Mode Transition Tool>transition cbt volumepair add -p sample_project
-v vol1 -c vol1 -g aggr1
```

5. Add the IP addresses to be transitioned to the project:

```
transition cbt lif add -p project_name -i ip-address [-m netmask] [-g default-gateway] [-p home_port] [ -n home_node]
```

To a transition project, you can also add existing IP addresses that are unconfigured or removed from the 7-Mode system during storage cutover, and new IP addresses that are not yet configured on either the 7-Mode or ONTAP.



FC and iSCSI LIFs are not transitioned by the 7-Mode Transition Tool.

```
7-Mode Transition Tool>transition cbt lif add -p sample_project -i
192.0.2.250 -m 255.255.255.128 -g 192.40.0.1 -p e0a -n cluster1-01
```

The new IP addresses are configured on the ONTAP system in the up state during the apply configuration (precutover) operation.

6. Follow these steps to add information to the primary and secondary project:

| If you want to add information about the... | Enter... |
|---|---|
| Primary 7-Mode system to the secondary project | <pre>transition cbt add-primary-seven-mode-system -p <i>project_name</i> -h <i>source-host</i> -f <i>management-ipaddress</i> -d <i>data-copy-ipaddress</i> [-m <i>multipathing-ipaddress</i>]</pre> <p>For more information, see CBT : How to provide all the required primary 7-Mode details for a given secondary project.</p> |
| Secondary cluster or SVM to the primary project | <pre>transition cbt add-secondary-cluster-mode-system -p <i>project_name</i> -h <i>c-mode-host-name</i> -v <i>vserver-name</i></pre> <p>For more information, see the CBT: How to add the required secondary cluster details to the primary project</p> |

project_name is the name of the transition project.

source-host is the primary 7-Mode storage system host name or IP address as seen in the `snapmirror status` command output of the secondary 7-Mode system.

For considerations about providing the 7-Mode primary system details, see the man pages.

management-ipaddress is the management IP address of the source host.

data-copy-ipaddress is the IP address over which data is copied.

multipathing-ipaddress is the additional IP address that is used for data copy.

c-mode-host-name is the FQDN or IP address of the cluster to which the secondary volumes in the

project have been transitioned.

vserver-name is the name of the SVM hosting the secondary volumes.

7. Create a data copy schedule:

```
transition cbt schedule add -p project_name -n schedule_name -d days-range -b start-time -e duration -u update-frequency -t available-transfers-percentage -c max-cap-concurrent-transfers -x project-snapmirror-throttle
```

The following command shows how to add a schedule that uses 100% of available concurrent SnapMirror transfers. However, it does not exceed the 25 concurrent SnapMirror transfers at any point of time.

```
transition schedule add -p sample_project -n dr_active -d 1-5 -b 23:30 -e 03:00 -c 25 -x 200 -u 00:30
```

8. View the detailed information about the transition project created:

```
transition cbt show -p project-name
```

Customizing the transition of 7-Mode configurations by using the CLI

By default, all 7-Mode configurations are transitioned to ONTAP. You can choose to exclude some or all the volume, NFS, CIFS, SAN, and name services configurations from transition by using the 7-Mode Transition Tool CLI. You can also choose to consolidate the 7-Mode NFS export rules and Snapshot schedules, and reuse an existing NFS export policy and Snapshot policy on the target SVM.

You must perform this task before the configuration is applied, after which any modification will be ignored.

The 7-Mode Transition Tool does not perform prechecks for the configuration that is excluded.

By default, all 7-Mode configurations are selected for transition.

It is a best practice to run the prechecks with all configurations first, and then exclude one or more configurations in the subsequent run of the prechecks. This helps you to understand which configurations are excluded from transition and which prechecks are skipped subsequently.

Steps

- Exclude and verify the configurations:
 - a. Exclude the configurations:

```
transition cbt property-set -p project_name -n config_property_name -v true
```

config_property_name is the configuration that you want to exclude.

[Configurations that can be excluded from transition](#)

- b. Verify the value of the property that is set for excluding the configuration:

```
transition cbt property-get -p project_name -n config_property_name
```

- Consolidate NFS export rules for transition:

- Consolidate similar 7-Mode NFS export rules to a single export policy in clustered Data ONTAP, which can then be applied to the transitioned volume or qtree:

```
transition cbt property-set -p project_name -n nfs-consolidate-similar-7mode-exports -v true
```

If the `nfs-consolidate-similar-7mode-exports` property is set to `false`, the 7-Mode Transition Tool creates a new NFS export policy in ONTAP for each 7-Mode NFS export rule.

- Reuse an existing NFS export policy on the SVM that matches the export policy that will be created by the tool, which can be applied to the transitioned volumes or qtrees:

```
transition cbt property-set -p project_name -n nfs-reuse-matching-svm-export-policies -v true
```

- Consolidate Snapshot schedules for transition:

- Consolidate similar 7-Mode Snapshot schedules to a single Snapshot policy in ONTAP, which can then be applied to the transitioned volume:

```
transition cbt property-set -p project_name -n consolidate-similar-7mode-snapshot-policies -v true
```

If the `consolidate-similar-7mode-snapshot-policies` property is set to `false`, the 7-Mode Transition Tool creates a new Snapshot policy in ONTAP for each Snapshot schedule.

- Reuse an existing Snapshot policy on the SVM that matches the Snapshot policy that will be created by the tool, which can be applied to the transitioned volumes:

```
transition cbt property-set -p project_name -n reuse-matching-svm-snapshot-policies -v true
```

Configurations that can be excluded from transition

You can customize the configuration transition by excluding some volume-level or SVM-level configurations for NFS, CIFS, SAN, and name services configurations from transition by specifying the property name with the `transition cbt property-set` command of the 7-Mode Transition Tool CLI.

NFS

| 7-Mode configuration to exclude | Property name |
|---------------------------------|---|
| Export policies | <code>ignore-nfs-exports-transition</code> |
| NFS options | <code>ignore-nfs-options-transition</code> |
| All NFS configurations | <code>ignore-all-nfs-configurations-transition</code> |

CIFS

| 7-Mode configuration to exclude | Property name |
|---------------------------------|---|
| Local users and groups | <code>ignore-local-users-groups-transition</code> |

| 7-Mode configuration to exclude | Property name |
|----------------------------------|--|
| Home directory paths | ignore-cifs-home-directory-paths-transition |
| Symbolic links | ignore-cifs-symlinks-transition |
| Widelinks | ignore-cifs-widelinks-transition |
| Shares and Share ACLs | ignore-cifs-shares-and-acls-transition |
| CIFS options | ignore-cifs-options-transition |
| Name mapping | ignore-cifs-name-mapping-transition |
| Audit configuration | ignore-cifs-audit-transition |
| Preferred domain controller list | ignore-cifs-preferred-domain-controllers-list-transition |
| All CIFS configurations | ignore-all-cifs-configurations-transition |

Name services

| 7-Mode configuration to exclude | Property name |
|---------------------------------|-------------------------------------|
| Netgroups | ignore-netgroups-transition |
| UNIX users and groups | ignore-unix-users-groups-transition |
| NIS | ignore-nis-transition |
| DNS | ignore-dns-transition |
| LDAP | ignore-ldap-transition |
| /etc/nsswitch.conf file | ignore-nsswitch-transition |
| LDAP-based user mapping | ignore-nmswitch-transition |
| /etc/hosts files | ignore-etc-hosts-transition |

| 7-Mode configuration to exclude | Property name |
|----------------------------------|---|
| All name services configurations | ignore-all-nameservices-configurations-transition |

SAN

| 7-Mode configuration to exclude | Property name |
|---------------------------------|--|
| igroup and LUN mapping | ignore-igroup-and-lunmapping-transition |
| All configurations | ignore-all-san-configurations-transition |

Snapshot schedules

| 7-Mode configuration to exclude | Property name |
|---------------------------------|---|
| Snapshot schedules | ignore-snapshot-schedule-transition |
| |  <p>If this option is set to true, the 'default' Snapshot policy is applied to the transitioned volumes.</p> |

Related information

[Supported and unsupported CIFS configurations for transition to ONTAP](#)

[NFS transition: supported and unsupported configurations, and required manual steps](#)

[Name services transition: supported and unsupported configurations, and required manual steps](#)

[SAN transition: supported and unsupported configurations, and required manual steps](#)

Examples of consolidating NFS export rules and Snapshot schedules for transition

You might want to review examples of how similar 7-Mode export rules and 7-Mode Snapshot schedules are consolidated to a single NFS export policy and a single Snapshot policy in ONTAP. You might also want to understand how the consolidated policies are assigned to the transitioned volumes or qtrees with or without reusing a matching existing policy on the target SVM.

Example of consolidating NFS export rules for transition

NFS export rules in 7-Mode and ONTAP before transition

7-Mode export rules

```

/vol/vol1      -sec=sys,rw,nosuid
/vol/vol2      -sec=sys,rw,nosuid
/vol/vol3      -sec=sys,rw,nosuid

```

Export policies existing in ONTAP

```

cluster-2::> vserver export-policy show -vserver vs1
Vserver          Policy Name
-----
vs1              default
vs1              export_policy_1

```

The existing export policy export_policy_1 has the following export rule:

```

cluster-2::> vserver export-policy rule show -vserver vs1 -policyname
export_policy_1

```

| Vserver | Policy Name | Rule Index | Access Protocol | Client Match | RO Rule |
|---------|-----------------|------------|-----------------|--------------|---------|
| vs1 | export_policy_1 | 1 | nfs | 0.0.0.0/0 | sys |

Export policies in ONTAP after transition with consolidation (no reuse)

Volumes vol1, vol2, and vol3 have similar export rules in 7-Mode; therefore, a new consolidated export policy, transition_export_policy_1, is assigned to these volumes after transition:

```

cluster-2::> vserver export-policy show -vserver vs1
Vserver          Policy Name
-----
vs1              default
vs1              export_policy_1
vs1              transition_export_policy_1
3 entries were displayed.

```

```
cluster-2::> vserver export-policy rule show -vserver vs1 -policyname
transition_export_policy_1
```

| | Policy | Rule | Access | Client | RO |
|---------|----------------------------|-------|----------|-----------|------|
| Vserver | Name | Index | Protocol | Match | Rule |
| vs1 | transition_export_policy_1 | 1 | nfs | 0.0.0.0/0 | sys |

```
cluster-2::> volume show -vserver vs1 -volume vol1,vol2,vol3 -fields
policy
```

| vserver | volume | policy |
|---------|--------|----------------------------|
| vs1 | vol1 | transition_export_policy_1 |
| vs1 | vol2 | transition_export_policy_1 |
| vs1 | vol3 | transition_export_policy_1 |

3 entries were displayed.

Export policies in ONTAP after transition with consolidation and reuse

Volumes vol1, vol2, and vol3 have similar export rules in 7-Mode; therefore, a consolidated export policy is assigned to these volumes after transition. The export policy, export_policy_1, which matches the 7-Mode export rules, already exists on the SVM. Therefore, the policy is applied to these volumes:

```
cluster-2::> vserver export-policy show -vserver vs1
```

| Vserver | Policy Name |
|---------|-----------------|
| vs1 | default |
| vs1 | export_policy_1 |

2 entries were displayed.

```
cluster-2::> vserver export-policy rule show -vserver vs1 -policyname
export_policy_1
```

| | Policy | Rule | Access | Client | RO |
|---------|-----------------|-------|----------|-----------|------|
| Vserver | Name | Index | Protocol | Match | Rule |
| vs1 | export_policy_1 | 1 | nfs | 0.0.0.0/0 | sys |

```
cluster-2::> volume show -vserver vs1 -volume vol1,vol2,vol3 -fields
policy
vserver volume policy
-----
vs1      vol1      export_policy_1
vs1      vol2      export_policy_1
vs1      vol3      export_policy_1
3 entries were displayed.
```

Example of consolidating Snapshot policies for transition

Snapshot schedules in 7-Mode and ONTAP before transition

7-Mode schedule

| 7-Mode volume | 7-Mode Snapshot schedule |
|---------------|--|
| vol1 | 0 2 4@8,12,16,20 (weekly Snapshot copies: 0, daily Snapshot copies: 2, hourly Snapshot copies: 6 at 2, 4, 8, 12, 16, 20 hours) |
| vol2 | 0 2 4@8,12,16,20 |
| vol3 | 0 2 4@8,12,16,20 |
| vol4 | 1 2 3@8,12,16 (weekly Snapshot copies: 1, daily Snapshot copies: 2, hourly Snapshot copies: 3 at 8,12,16 hours) |
| vol5 | 2 2 3@8,12,16 (weekly Snapshot copies: 2, daily Snapshot copies: 2, hourly Snapshot copies: 3 at 8,12,16 hours) |

Snapshot policies existing in ONTAP

| Snapshot policy name | Policy details |
|----------------------|---|
| ScheduleWeekly | Weekly, count: 1 |
| ScheduleDailyHourly4 | Schedule details <ul style="list-style-type: none"> Schedule1: daily, count1: 2 Schedule2: hourly, count2: 4 every 8, 12, 16, 20 hours |
| ScheduleHourly1 | Hourly at 8, 12, 16, 20 hours, count: 4 |

Snapshot policy in ONTAP after transition with consolidation (no reuse)

| 7-Mode volume | 7-Mode Snapshot schedule | Snapshot policy in ONTAP |
|---|--|---|
| vol1 | 0 2 4@8,12,16,20 (weekly Snapshot copies: 0, daily Snapshot copies: 2, hourly Snapshot copies: 4 at 8, 12, 16, 20 hours) | Consolidated policy for vol1, vol2, and vol3 <ul style="list-style-type: none"> • Name: transition_snapshot_policy_0 • Schedule details <ul style="list-style-type: none"> ◦ Schedule1: daily, count1: 2 ◦ Schedule2: hourly, count2: 4 every 8, 12, 16, 20 hours |
| vol2 | 0 2 4@8,12,16,20 | vol3 |
| 0 2 4@8,12,16,20 | vol4 | 1 2 3@8,12,16 (weekly Snapshot copies: 1, daily Snapshot copies: 2, hourly Snapshot copies: 3 at 8,12,16 hours) |
| <ul style="list-style-type: none"> • Name: transition_snapshot_policy_1 • Schedule details <ul style="list-style-type: none"> ◦ Schedule1: weekly, count1: 1 ◦ Schedule2: daily, count2: 2 ◦ Schedule3: hourly, count3: 3 every 8,12,16 hours | vol5 | 2 2 3@8,12,16 (weekly Snapshot copies: 2, daily Snapshot copies: 2, hourly Snapshot copies: 3 at 8,12,16 hours) |

Snapshot policy in ONTAP after transition with consolidation and reuse

| 7-Mode volume | 7-Mode Snapshot schedule | Snapshot policy in ONTAP |
|------------------|--|--|
| vol1 | 0 2 4@8,12,16,20 (weekly Snapshot copies: 0, daily Snapshot copies: 2, hourly Snapshot copies: 4 at 2, 4, 8, 12, 16, 20 hours) | Consolidated policy for vol1, vol2, and vol3 for which the existing ONTAP policy is reused Name: ScheduleDailyHourly4 |
| vol2 | 0 2 4@8,12,16,20 | vol3 |
| 0 2 4@8,12,16,20 | vol4 | 1 2 3@8,12,16 (weekly Snapshot copies: 1, daily Snapshot copies: 2, hourly Snapshot copies: 3 at 8,12,16 hours) |

| 7-Mode volume | 7-Mode Snapshot schedule | Snapshot policy in ONTAP |
|---|--------------------------|---|
| <ul style="list-style-type: none"> • Name: transition_snapshot_policy_1 • Schedule details <ul style="list-style-type: none"> ◦ Schedule1: weekly, count1: 1 ◦ Schedule2: daily, count2: 2 ◦ Schedule3: hourly, count3: 3 every 8,12,16 hours | vol5 | 2 2 3@8,12,16 (weekly Snapshot copies: 2, daily Snapshot copies: 2, hourly Snapshot copies: 3 at 8,12,16 hours) |

Running prechecks

After creating a transition session, you can validate the session to verify whether the 7-Mode environment can be transitioned by using the `transition precheck` command. You can verify the compatibility between the 7-Mode storage system and the SVM for factors such as features and security settings.

The transition session must be in the preparation state.

The `transition cbt precheck` command generates a detailed error message report. Issues identified in the report are assigned different severity levels, depending on the impact of the issue on the transition process. Because some errors can hinder the transition, you must perform corrective actions for errors and warnings. You should also review the impact of messages of other severity levels before proceeding with the transition.

Steps

1. Verify the compatibility for transition:

```
transition cbt precheck -p project_name
```



```

7-Mode Transition Tool>transition cbt precheck -p project_finance

[1/14 ] Project status checks
[   Errors   ]

Operation summary:
-----
The 'precheck' operation on the project 'project_finance' has failed.

1  Errors - Failed!

Do you want to see a detailed report {yes, no} [yes]:

1 Errors:
-----
90202: Checking whether any of the 7-Mode volumes are in 'offline'
state.                                [   Error   ]
  > Following volumes are in 'offline' state.
  > CORRECTIVE-ACTION: Bring offline volumes online by using the
following command: 'vol online
  > <volume_name>'.
  > AFFECTED-OBJECTS: { vol2 }

Next suggested steps:
-----
1. Review the 'precheck' operation results, fix all blocking errors and
run the 'precheck' operation again.
2. Use the 'transition job-results -j 6e33e0a7-bb36-49df-91f3-
2e52cbfa3074' command to see the results of this
operation.

Ran precheck for project 'project_finance'.

```

You must resolve all of the errors that might cause problems with the transition before starting the data copy from the 7-Mode volumes.

Starting data copy for a transition project

You can initiate a data copy from 7-Mode volumes to ONTAP volumes for a baseline transfer. The baseline transfer starts when the data copy schedules configured in a project becomes active. After the baseline transfer is complete, the ONTAP volumes are updated periodically based on the data copy schedule.

- You must have run a precheck on the project and resolved all errors.
- You must have configured a dedicated, high-bandwidth, low-latency network between the source and

destination systems for data transfer.



You can use SnapMirror multipathing to balance the data transfer load between the two paths.

- A data copy schedule must be configured for the session.

If you are transitioning to an existing SVM, the volumes are created during this operation. If you are transitioning to a new SVM, the SVM and its volumes get created during this operation.

SnapMirror relationships are created between 7-Mode and ONTAP volumes.

Steps

1. Start the data copy:

transition start -p *project_name*

If the start operation fails and data copy is not started, you must resolve the issues and run the `transition start` command again.

```
7-Mode Transition Tool>transition start -p project_finance

[1/17 ] Project status checks                               [
Ok      ]

[2/17 ] Validating 7-Mode system information                 [
Ok      ]

[3/17 ] Validating 7-Mode volumes information                [
Ok      ]

[4/17 ] Validating Cluster-Mode system information           [
Ok      ]

[5/17 ] Validating Cluster-Mode volume information           [
Ok      ]

[6/17 ] Validating LIF information                           [
Notifications ]

[7/17 ] WAFL prechecks                                       [
Ok      ]

[8/17 ] UNIX users and groups prechecks                     [
Warnings  ]

[9/17 ] SnapMirror prechecks                                 [
Notifications ]

[10/17] NFS prechecks                                        [
Warnings  ]

[11/17] Networking prechecks                                 [
Notifications ]

[12/17] CIFS prechecks                                       [
Warnings  ]

[13/17] Name services prechecks                              [
Warnings  ]
```

```

[14/17] SAN prechecks
Ok
[15/17] Creating Cluster-Mode volumes
Ok
[16/17] Establishing SnapMirror relationships between the 7-Mode and
Cluster-Mode volumes
[      Ok      ]
[17/17] Initializing SnapMirror relationships between the 7-Mode and
Cluster-Mode volumes
[      Ok      ]

Operation summary:
-----
The 'start' operation is completed with warnings and notifications.

0  Errors - No action required
22  Warnings - Need your attention!
16  Notifications - Please review
102  Informational messages

Do you want to see a detailed report {yes, no} [yes]:

22 Warnings:
-----
20362: Checking whether 7-Mode UNIX groups' names are in use in the
Vserver 'vs2'. [      Warning      ]
> The following 7-Mode UNIX groups' names are already in use in the
Vserver 'vs2'.
> AFFECTED-OBJECTS: { daemon }

20372: Checking whether 7-Mode UNIX groups' ids are in use in the
Vserver 'vs2'. [      Warning      ]
> The following 7-Mode UNIX groups' ids are already in use in the
Vserver 'vs2'.
> AFFECTED-OBJECTS: { daemon }

```

The project moves to the copy-baseline state. After the baseline transfer is completed, the project moves to the copy-update state.

Applying configurations to ONTAP volumes

After the baseline data copy is completed, you can copy and apply all NAS configurations from the 7-Mode system (including protocols and services configuration) to the ONTAP volumes. If the target cluster is running Data ONTAP 8.3.2 or later, SAN configurations are transitioned in this phase.

- The configurations are applied in the precutover phase, which has two modes: precutover read-only mode and precutover read/write mode.
- The precutover read/write mode is not supported when the project contains the following:
 - SAN volumes and the target cluster is running Data ONTAP 8.3.1 or earlier. In this situation, the following configurations are not applied in the precutover phase, instead they are applied during the cutover phase:
 - SAN configurations
 - Snapshot Schedule configurations
 - SnapLock Compliance volumes.

If the project contains SnapLock Compliance volumes, then the Snapshot Schedule configurations are not applied in the precutover phase, instead these configurations are applied during the cutover phase.

See [Considerations for transitioning of SnapLock Compliance volumes](#).

- If the target cluster is running Data ONTAP 8.3.1 or earlier, the read/write mode is not supported during precutover if any volume in the project contains LUNs.
- If the target cluster is running Data ONTAP 8.3.1 or earlier and a project contains both SAN and NAS volumes, only NAS configurations are transitioned in the precutover phase, and these NAS configurations can be tested only in the read-only mode.

SAN configurations are transitioned only in the storage cutover phase.



SAN LIFs are not configured.

- If the target cluster is running Data ONTAP 8.3.1 or earlier and the project contains SAN volumes, then Snapshot schedule configuration is not applied during apply configuration (precutover) phase.

These configurations are applied during the cutover phase.

- If the target cluster is running Data ONTAP 8.3.1 or earlier and the project contains SAN volumes, then Snapshot schedule configuration is not applied during apply configuration (precutover) phase.

These configurations are applied during the cutover phase.

Steps

- Run the transition operation in read-only mode:

```
transition cbt precutover -p project_name -m ro_test -c  
max_cap_concurrent_transfers -b snapmirror_throttle
```

- Run the transition operation in read/write mode:

- Apply the configurations to ONTAP: **+ transition cbt precutover -p *project_name* -m *rw_test* -c *max_cap_concurrent_transfers* -b *snapmirror_throttle***
- Manually test and verify the configurations in ONTAP.
- Start the resynchronization operation: **+ transition cbt resync -p *project_name***

Completing the transition

Because the complete operation is disruptive, you should evaluate when to run it. When transitioning volumes in a SnapMirror relationship, the secondary project must be completed before completing the transition of the primary project.

Steps

1. Perform a data copy from the 7-Mode volumes to the ONTAP volumes:

```
transition cbt update -p project_name -t available-transfers-percentage -c snapmirror-throttle max-cap-concurrent-transfers -x -r interactive
```

This helps in reducing the downtime during storage cutover.

2. Manual: Disconnect client access.
3. Complete the transition process:

```
transition cbt cutover -p project_name -t available-transfers-percentage -c max-cap-concurrent-transfers -x snapmirror-throttle -r interactive -o [true|false]
```

By default, the 7-Mode volumes are made offline. To keep the 7-Mode volumes online, set the `-o` option to `false`.



If the cutover is not successful for some reason—for example, the 7-Mode system or cluster is not reachable—you must resolve the issues and run the `transition cutover` command again.

For more information about the commands, see the man pages.

The transition session moves to the cutover state. If the cutover is successful, the session moves to the completed state.

You must perform any post-transition tasks and reconnect the clients to the ONTAP volumes.

Commands for managing transition

You can use the `transition cbt` command to manage functions, such as creating, modifying, deleting, and displaying transition projects.

The following option can be used with all commands:

`-r no`: Disables the interactive nature of the command. By default, this option is enabled.

You can also use the `-v` option, which enables the verbose mode, with the `transition version` and `transition cbt menu` commands.

Commands to manage transition operations

| If you want to... | Use this command... |
|--|--|
| Create a transition project | <code>transition cbt create</code> |
| Add the cluster information to the given project | <code>transition cbt add-cluster-info</code> |
| Modify a transition project | <code>transition cbt modify</code> |
| Complete the transition to an ONTAP storage system | <code>transition cbt cutover</code> |
| Run precheck on the project | <code>transition cbt precheck</code> |
| Apply the configurations to an ONTAP storage system | <code>transition cbt precutover</code> |
| Trigger the SnapMirror resynchronization between the 7-Mode volumes and ONTAP volumes in the project | <code>transition cbt resync</code> |
| Trigger the SnapMirror update operation for the ONTAP volumes | <code>transition cbt update</code> |
| Display the version of 7-Mode Transition Tool | <code>transition cbt version</code> |
| Start baseline data copy from the 7-Mode volumes to ONTAP volumes in a project | <code>transition cbt start</code> |
| Pause data copy from the 7-Mode volumes to ONTAP volumes in a project | <code>transition cbt pause</code> |
| Resume data copy transfers from the 7-Mode volumes to ONTAP volumes of a paused project | <code>transition cbt resume</code> |
| Abort a transition project | <code>transition cbt abort</code> |
| Delete a transition project | <code>transition cbt delete</code> |
| Display the list of transition projects or display information about a project | <code>transition cbt show</code> |

Commands to manage volume SnapMirror relationships

| If you want to... | Use this command... |
|---|---|
| Add information about a 7-Mode primary system to a secondary project | <pre>transition cbt add-primary-seven-mode-system `</pre> <p>CBT : How to provide all the required primary 7-Mode details for a given secondary project</p> |
| Add information about an ONTAP secondary system to a primary project | <pre>transition cbt add-secondary-cluster-mode-system</pre> <p>CBT: How to add the required secondary cluster details to the primary project</p> |
| Remove information about a 7-Mode primary system from a secondary project | <pre>transition cbt remove-primary-seven-mode-system</pre> |
| Remove information about an ONTAP secondary system from a primary project | <pre>transition cbt remove-secondary-cluster-mode-system</pre> |
| List the 7-Mode primary systems added to a secondary project | <pre>transition cbt show-primary-seven-mode-systems</pre> |
| List the ONTAP secondary systems added to a primary project | <pre>transition cbt show-secondary-cluster-mode-systems</pre> |

Commands to manage credentials of Data ONTAP systems

| If you want to... | Use this command... |
|--|--|
| Save credentials (user name and password) of a host in the transition server | <pre>transition credentials add</pre> |
| Retrieve user name of a host | <pre>transition credentials get</pre> |
| Modify the user name or password of a host | <pre>transition credentials modify</pre> |
| Delete credentials of a host from the transition server | <pre>transition credentials remove</pre> |

Commands to manage the volume pairs in a project

| If you want to... | Use this command... |
|---|--|
| Add the 7-Mode volume and the ONTAP volume (volume pair) to a project | <pre>transition cbt volumepair add</pre> |

| If you want to... | Use this command... |
|--------------------------------------|---|
| Modify the ONTAP volume of a project | <code>transition cbt volumepair modify</code> |
| Remove a volume pair from a project | <code>transition cbt volumepair remove</code> |
| List volume pairs in a project | <code>transition cbt volumepair show</code> |

Commands to manage LIF details



FC and iSCSI LIFs are not transitioned by the 7-Mode Transition Tool.

| If you want to... | Use this command... |
|--------------------------------------|--|
| Add anSVM data LIF to a project | <code>transition cbt lif add</code> |
| Modify anSVM data LIF of a project | <code>transition cbt lif modify</code> |
| Remove LIFs from a project | <code>transition cbt lif remove</code> |
| List all the LIFs added to a project | <code>transition cbt lif show</code> |

Commands to manage project properties

| If you want to... | Use this command... |
|--|--|
| Set value for a specific project property | <code>transition cbt property-set</code> |
| Clear the value of a specific project property | <code>transition cbt property-reset</code> |
| Get the value of a specific project property | <code>transition cbt property-get</code> |

Commands to manage transition jobs

| If you want to... | Use this command... |
|--|-------------------------------------|
| List jobs that ran or are running on the given project and operation | <code>transition jobs</code> |
| View the status of a job | <code>transition job-status</code> |
| View the results of a job | <code>transition job-results</code> |

Commands to manage transition schedules

| If you want to... | Use this command... |
|--|---|
| Add a schedule to manage SnapMirror transfers along with bandwidth | <code>transition cbt schedule add</code> |
| Modify a SnapMirror schedule of the project | <code>transition cbt schedule modify</code> |
| Remove the SnapMirror schedules from the project | <code>transition cbt schedule remove</code> |
| List all the SnapMirror schedules in a project | <code>transition cbt schedule show</code> |

Command to collect tool logs

| If you want to... | Use this command... |
|--|--|
| Collect the log files of 7-Mode Transition Tool Logs are saved on the server in the <code>asup</code> directory of the 7-Mode Transition Tool installation path. | <code>transition bundle-tool-logs</code> |

For more information about these commands, see the man pages for the 7-Mode Transition Tool CLI.

Related information

[Migrating data and configuration from 7-Mode volumes](#)

Performing manual post-transition tasks

After the storage cutover finishes successfully and transition is completed, you must perform some post-transition tasks for configuring features that are not transitioned, before enabling client access to the ONTAP volumes.

Steps

1. Manual: Perform the required steps to configure features that were not transitioned, as listed in the precheck report.

For example, you must manually configure IPv6 and FPolicy after transition.

2. Delete or consolidate the configurations that are transitioned.

The 7-Mode Transition Tool creates a new instance for all configurations that are transitioned from 7-Mode. Therefore, there might be several instances of some configurations, such as cron schedules, that might need to be consolidated or deleted manually.

3. SAN transition: Manually reconfigure the hosts.

[SAN host transition and remediation](#)

4. Ensure that the SVM is ready to serve data to the clients by verifying the following:

- The volumes on the SVM are online and read/write.
 - The IP addresses on the SVM are up and reachable from the name servers.
5. Redirect client access to the ONTAP volumes.

Related information

[NetApp Documentation: ONTAP 9](#)

Configuring zones by using the FC zone plan

After completing the transition, you must configure the zones by using the FC zone plan to group the initiator hosts and targets for providing data access from the cluster.

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

Steps

1. Log in to the CLI of the switch.
2. 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 pwwn 21:00:00:c0:dd:19:4b:f3
member pwwn 20:07:00:a0:98:32:99:07
member pwwn 20:09:00:a0:98:32:99:07
.....
.....
.....
copy running-config startup-config
```

The initiator hosts can access data from the cluster.

Recovering from a failed LUN transition

If the transition of volumes with LUNs fails, you can use the `lun transition 7-mode show` command to check which LUNs were not transitioned to ONTAP, and then determine a corrective action.

Steps

1. Change to advanced privilege level:

```
set -privilege advanced
```

2. Check which LUNs failed:

```
lun transition 7-mode show
```

3. Review the EMS logs and determine the corrective action that you must take.

4. Perform the required steps shown in the EMS message to correct the failure.

5. If any supported LUNs failed the transition, then to complete the transition:

```
lun transition start
```

6. View the transition status of the volumes:

```
lun transition show
```

The transition status can be one of following values:

- **active**: The volume is in an active SnapMirror transition relationship and not yet transitioned.
- **complete**: All supported LUNs are transitioned for this volume.
- **failed**: LUN transition failed for the volume.
- **none**: The volume did not contain LUNs to transition from 7-Mode systems.

```
cluster1::*> lun transition show
Vserver          Volume          Transition Status
-----
vs1              vol0            none
                 vol1            complete
                 vol2            failed
                 vol3            active
```

Viewing transitioned SAN configurations

The LUN serial numbers and LUN IDs of SAN volumes change after transition. To view the transitioned SAN configurations, you can generate the mapping of the old LUNs to the new transitioned LUNs by using the 7-Mode Transition Tool CLI and export the LUN mapping information to a `.csv` file.

Storage cutover must have completed successfully.

Steps

1. To generate LUN mapping information, run the following command from the CLI:

```
transition cbt export lunmap -p project_name -o file_path
```

```
C:\Program Files\NetApp\7-Mode Transition Tool>transition cbt export  
lunmap -s SanWorkLoad -o c:/Libraires/Documents/7-to-C-LUN-MAPPING.csv  
Successfully exported 7 to C LUN mapping for 'SanWorkLoad'
```

The following example shows a sample LUN mapping file:

```
7-Mode Storage System,192.168.43.49,  
vfiler,finance,  
  
cDOT Storage System,192.168.32.97,  
SVM,finance,  
  
LUN#,Source LUN Path,7-Mode Serial Number,Source Igroups,Source mapping  
LUN ID,Destination LUN Path,Serial Number,Destination Igroup,Destination  
mapping LUN ID  
LUN#1,/vol/vol_SANdataset_sm_0/lun-  
inventory,dtY5B4tFAaAF,WinGrp,0,/vol/vol_SANdataset_sm_0/lun-  
inventory,7SQ8p$DQ12rX,WinGrp,0  
LUN#1,/vol/vol_SANdataset_sm_0/lun-  
inventory,dtY5B4tFAaAF,WinGrp2,3,/vol/vol_SANdataset_sm_0/lun-  
inventory,7SQ8p$DQ12rX,WinGrp2,3  
LUN#1,/vol/vol_SANdataset_sm_0/lun-  
inventory,dtY5B4tFAaAF,WinGrp3,4,/vol/vol_SANdataset_sm_0/lun-  
inventory,7SQ8p$DQ12rX,WinGrp3,4  
LUN#2,/vol/vol_SANdataset_sm_0/lun-  
payroll,dtY5B4tFAaAC,LnxGrp1,2,/vol/vol_SANdataset_sm_0/lun-  
payroll,7SQ8p$DQ12rT,LnxGrp1,4  
LUN#2,/vol/vol_SANdataset_sm_0/lun-  
payroll,dtY5B4tFAaAC,LnxGrp2,2,/vol/vol_SANdataset_sm_0/lun-  
payroll,7SQ8p$DQ12rT,LnxGrp2,4
```

You must perform the required post-transition tasks on the host before restoring access to the transitioned ONTAP volumes.

[SAN host transition and remediation](#)

Limitations with 7-Mode Snapshot copies of LUNs managed by SnapDrive and SnapManager after transition

When transitioning to clustered Data ONTAP 8.3, SnapDrive and SnapManager backup verification and restore operations fail on the transitioned 7-Mode Snapshot copies. When transitioning to clustered Data ONTAP 8.3.1, SnapManager for Microsoft Exchange Server (SME) backup verification and SnapManager for Hyper-V (SMHV) restore operations of files fail on the transitioned 7-Mode Snapshot copies.

Depending on the SnapDrive or SnapManager that is used on the host connected to the 7-Mode system and the Data ONTAP version running on your cluster, 7-Mode Transition Tool displays a warning message for this limitation during precheck.



All SnapDrive and SnapManager backup verification and restore operations are supported on any newly created Snapshot copies in clustered Data ONTAP.

Workaround for backup and restore operations using SnapDrive and SnapManager in clustered Data ONTAP 8.3

Perform the one of the following actions:

- Retain the 7-Mode systems till the 7-Mode Snapshot copies expire.

You can use the 7-Mode system to restore a file from a LUN in the 7-Mode Snapshot copy and then migrate it to clustered Data ONTAP, if required.

- Use single file SnapRestore in ONTAP to restore the LUN from a 7-Mode Snapshot copy to the active file system, and then mount the restored LUN for single file restore.

Workaround for backup and restore operations using SME and SMHV in ONTAP 8.3.1 and later supported releases

Perform the one of the following actions:

- Verify that the SnapManager backups created in 7-Mode are valid before you perform a storage cutover to ONTAP.
- Create a clone of the volume in ONTAP from the transitioned 7-Mode Snapshot copies, and then mount the LUN on the cloned volume for backup verification and restore operations.

The best practice is to retain the 7-Mode systems till the 7-Mode Snapshot copies expire.

[NetApp Documentation: SnapManager for Microsoft Exchange Server](#)

[NetApp Documentation: SnapManager for Hyper-V](#)

[NetApp Documentation: SnapManager for SAP](#)

[NetApp Documentation: SnapManager for Oracle](#)

[NetApp Documentation: SnapDrive for UNIX](#)

[NetApp Documentation: SnapDrive for Windows \(current releases\)](#)

[NetApp Documentation: SnapManager for Microsoft SQL Server](#)

[NetApp Documentation: SnapManager for Microsoft SharePoint](#)

Configuring cron job schedules on the remote site after transitioning a MetroCluster configuration

The cron job schedules created during transition are not replicated to the remote site, and therefore negotiated switchover fails after transition. You must manually create the cron job schedules on the remote site after transition.

Storage cutover for the 7-Mode active site must have been completed successfully.

Steps

1. Record the cron job messages generated at the storage cutover to identify the job schedules that must be replicated to the remote site.

| ID | Type(Tags) | Message | Recommended Action | Affected Objects |
|-------|-------------|---|---|------------------|
| 88101 | Warning | Cron jobs have been configured by the 7-Mode Transition Tool on the local site of this MetroCluster configuration. These cron jobs on the local cluster must be applied to the remote cluster of this MetroCluster configuration to support negotiated switchover and switchback operations. | Execute the following cluster commands on the remote site of this MetroCluster configuration. | Affected Objects |
| 90702 | Warning | Affected Objects : MetroCluster Cron Jobs # Cron job schedules # For Transition to MetroCluster, the following schedules must be applied to the remote site to enable negotiated switchover and switchback. job schedule cron create -name CronJob_mcc_vserver1_mcc_volt_0 -hour 0 -minute 0 job schedule cron create -name CronJob_mcc_vserver1_mcc_volt_1 -hour 8,12,16,20 -minute 0 | | |
| 91061 | Information | | | |

2. From the cluster CLI of the remote site, run the `job schedule cron create` commands recorded in [Step 1](#).

Deleting 7-Mode volume Snapshot copies from transitioned volumes

Because the 7-Mode volume Snapshot copies are not deleted automatically from a transitioned volume, you must delete these Snapshot copies manually after the transition operation is completed.

Because the naming convention for Snapshot copies differs between 7-Mode environments and clustered Data ONTAP, the 7-Mode Snapshot copies cannot be automatically overwritten by the Snapshot copies in the clustered Data ONTAP environment.

To achieve better space utilization, you should delete the unwanted 7-Mode Snapshot copies, including the ones from the 7-Mode to 7-Mode SnapMirror relationships and the 7-Mode to clustered Data ONTAP SnapMirror relationships.

Steps

1. After the storage cutover operation is completed, delete the unwanted 7-Mode Snapshot copies:

```
snap delete -vserver <svm_name> -volume <volume_name> -snapshot  
<snapshot_name>
```

Consolidating cron schedules from transitioned volumes

The 7-Mode Transition Tool generates unique schedules per transitioned volume. Because of this, a large number of cron schedules can be created after the transition, so you need to know how to consolidate them after transition.

Steps

1. Consolidate the cron schedules after transition by using the `job schedule cron` command.

Managing a transition project

You can manage transition projects by using the 7-Mode Transition Tool (web interface). You can edit, pause, resume, abort, delete, and run prechecks on a project.

Editing a project

You can edit a project to add or remove volumes, and to modify the transition configurations.

- You must have created a project.
- You must not have started the project that you want to edit.

You can edit a project to change the volumes that are selected for transition. You can also edit a project to change the SVM mapping, volume mapping, and interface mapping, as well as the data copy schedule. You cannot change the project type or the vFiler unit when editing a project.

Steps

1. From the **Dashboard** window, select the project that you want to edit.
2. Click **Edit Project**, and then select the required step from the drop-down list.
3. Complete the wizard to edit your project.

Managing SnapMirror transfers and schedule

In the data copy phase, the clustered Data ONTAP volumes are updated periodically with the data from the 7-Mode volumes based on a data copy schedule. You can create, edit, and delete the data copy schedule to be used for SnapMirror baseline copy, update, and resynchronization operations.

You can manage the copy operations for all the volumes selected for transition by specifying multiple data copy schedules comprising details such as, start time, duration, update frequency, the number of concurrent SnapMirror transfers to run during that schedule, and throttle limit.

You can provide multiple concurrent transfers count and throttle limit for different time periods, such as weekday, weekend, business hours, non-business hours, DR schedules, and non-DR schedules.

Related information

[Creating a transition project](#)

Creating a data copy schedule for SnapMirror transfers

You can create data copy schedules to effectively manage the transition data copy operations for the volumes in a project. You can also specify the number of concurrent SnapMirror transfers to run during that schedule to ensure that any replication operation does not fail due to the concurrent SnapMirror transfers reaching the maximum limit.

- The project must be in the preparation, data copy, or apply configuration (precutover) phase.
- The 7-Mode Transition Tool service must always be running for the schedules to be effective.



The data copy schedule is used for SnapMirror baseline copy, update, and resynchronization operations.

- You must create at least one data copy schedule for each project.
- You can create a maximum of 7 schedules per project; however, the schedules cannot overlap within a project.

For example, you can create customized schedules for business hours and non-business hours, DR hours and non-DR hours, and weekdays and weekends.



If the projects use the same 7-Mode controller or cluster, ensure that the data copy schedules do not overlap across different projects.

- The configured schedules are applied based on the 7-Mode controller's time zone.
- The number of concurrent SnapMirror transfers to use during the schedule is determined at run time based on the percentage of concurrent SnapMirror streams to use and the maximum limit configured.
- The number of concurrent SnapMirror transfers to use during the schedule should be provided in such a way that the existing 7-Mode DR schedules are not affected because of the tool using the SnapMirror transfers specified with the schedule.
- If the number of concurrent SnapMirror transfers that the tool is using is less than the configured number of concurrent SnapMirror transfers to use during the schedule, tool schedules new transfers to make use of the remaining transfers.
- If the schedule is coming to an end or there is a reduction in number of concurrent SnapMirror transfers on the 7-Mode storage system, the tool aborts the extra transfers to ensure that it uses only the configured number of transfers at any point in time.



If a baseline is in progress and Snapshot checkpoint is not yet created, the tool does not abort the transfer but waits for the Snapshot checkpoint to be created before aborting the transfer.

Steps

1. Create schedules from the Data Copy Schedule page of the Dashboard by clicking **Edit Project**, and then selecting **Configure Schedule**.
2. Enter a name for the new schedule.
3. In the Recurring Days pane, select the days when the data copy operation should run.

| If you want to run the data copy operations... | Then... |
|--|---|
| Daily | Select Daily . This is the default value. |
| Only on certain days | <ol style="list-style-type: none">a. Select Select Days.b. Select the days of the week. |

4. In the Time Interval pane, specify the start time, duration, and frequency of the data copy schedule.
 - a. Enter the time the data copy operations should start by selecting the hours and minutes from the **Start Time** drop-down list.

Valid values are from 00:00 to 23:30.

- b. Enter the time period for which you want the data copy operations to run, select the hours and minutes from the **Duration** drop-down list.



The duration of a schedule should not exceed a week (167 hours and 30 minutes).

For example, if 20:30 is specified, SnapMirror operations run for next 20 hours and 30 minutes from the start time.

- c. Select the frequency at which incremental transfers should be done (within the created schedule duration) after the baseline transfer has been completed by doing one of the following:
 - Select the hours and minutes from the **Update Frequency** drop-down list.
 - Select **Continuous Updates**.

The minimum delay between two consecutive updates will be 5 minutes.

By default, SnapMirror updates run every 30 minutes.

5. In the Parameters for Transition Data Copy Operations pane, enter the SnapMirror parameters.

- a. Specify the maximum number of concurrent SnapMirror transfers that should be used for data copy by doing one or all of the following:
 - Specify the percentage of the available volume SnapMirror transfers that should be used for data copy (when schedule is active), by entering the percentage in the **Maximum Number of Concurrent VSM Transfers** field.

The available volume SnapMirror transfers is calculated at run time.



The maximum number of concurrent SnapMirror transfers supported on your platform is displayed in this pane.

- Specify the maximum number of concurrent volume SnapMirror transfers that can run during this schedule in the **Not Exceeding** field. If you enter both values, the lowest value is used as the number of concurrent transfers.

The number of concurrent transfers to be used for transition is calculated at run time based on the schedule and the number of configured concurrent transfers.

+ Your platform supports a maximum of 100 concurrent volume SnapMirror transfers, 60 are currently available, and you have specified the following values:

- Percentage of the available volume SnapMirror transfers option is 50%.

The maximum number of concurrent transfers based on the percentage option is 50% of 60 = 30.

- Maximum number of concurrent volume SnapMirror transfers option is 25. In this scenario, the tool sets the maximum number of concurrent volume SnapMirror transfers to 25, which is the lowest of the two values.

- a. Specify the maximum bandwidth in MB/s (throttle) by doing one of the following:

| If you want to... | Then... |
|-------------------------------------|--|
| Utilize all the available bandwidth | Select Maximum . This is the default value. |
| Specify the throttle value | Enter the value in the Not Exceeding field. Maximum allowed input value is 4194303. + |

The throttle value is equally distributed among all the active transfers in the project.



The throttle for each transfer is determined at run time based on the number of available concurrent volume SnapMirror transfers.

If the active schedule is configured with the throttle value of 200 MBps and only 10 concurrent transfers are available, each transfer uses 20 MBps bandwidth.

The schedules become effective only when the project is in data copy or apply configuration (precutover) phase.

Example for planning a data copy schedule

Consider a 7-Mode controller that supports 100 concurrent SnapMirror transfers with 75 DR relationships. The business requirements need SnapMirror operations to run during the following timings:

| Days | Time | Currently used SnapMirror transfers |
|--------------------|---|-------------------------------------|
| Monday to Friday | 9:00 a.m. to 5:00 p.m. | 50% of available transfers |
| Monday to Friday | 11:30 p.m. to 2:30 a.m. | 75 transfers used for DR |
| Monday to Friday | 2:30 a.m. to 9:00 a.m. and 5:00 p.m. to 11:30 p.m. | 25% of available transfers |
| Saturday to Monday | 2:30 a.m. (Saturday) to 9:00 a.m. (Monday) | 10% of available transfers |

You can create the following data copy schedules to manage your transition data copy operations:

| Schedule | Option | Value |
|------------|--|------------------|
| peak_hours | Days Range | Monday to Friday |
| | Start Time | 09:30 |
| | Duration | 8:00 |
| | Percentage of maximum number of concurrent transfers | 50 |
| | Maximum number of concurrent transfers | |
| | Throttle (MBps) | 100 |
| | Update Frequency | 0:00 |
| dr_active | Days Range | Monday to Friday |
| | Start Time | 23:30 |
| | Duration | 3:00 |
| | Percentage of maximum number of concurrent transfers | |
| | Maximum number of concurrent transfers | 25 |
| | Throttle (MBps) | 200 |
| | Update Frequency | 0:30 |

| Schedule | Option | Value |
|------------------|--|------------------|
| non_peak_non_dr1 | Days Range | Monday to Friday |
| | Start Time | 17:00 |
| | Duration | 6:30 |
| | Percentage of maximum number of concurrent transfers | 75 |
| | Maximum number of concurrent transfers | |
| | Throttle (MBps) | 300 |
| | Update Frequency | 1:00 |
| non_peak_non_dr2 | Days Range | Monday to Friday |
| | Start Time | 02:30 |
| | Duration | 6:30 |
| | Percentage of maximum number of concurrent transfers | 75 |
| | Maximum number of concurrent transfers | |
| | Throttle (MBps) | 300 |
| | Update Frequency | 1:00 |

| Schedule | Option | Value |
|-----------|--|----------|
| week_ends | Days Range | Saturday |
| | Start Time | 02:30 |
| | Duration | 53:30 |
| | Percentage of maximum number of concurrent transfers | 90 |
| | Maximum number of concurrent transfers | |
| | Throttle (MBps) | 500 |
| | Update Frequency | 2:00 |

Considerations for creating a data copy schedule

The 7-Mode Transition Tool runs a scheduler that checks for the active schedule every 5 minutes. You must be aware of the requirements to make a data copy schedule active. You can manage the SnapMirror transfers effectively by following some best practices when configuring the different parameters of a data copy schedule.

Requirements to make a data copy schedule active

- The 7-Mode Transition Tool service must be running.

If the 7-Mode Transition Tool service is restarted, the SnapMirror operations are not performed until the credentials are added.

- There should be at least one data copy schedule available for SnapMirror data copy operations to run.

If there is no schedule available for a particular time range, the SnapMirror data copy operations are not performed during that time.

- If the SnapMirror relationships are in the quiesced state, the data copy operations are not performed.
- The system time of the 7-Mode and the cluster must be synchronized for the incremental transfers to happen according to the data copy schedule.

If the 7-Mode system time is lagging behind the cluster time, the updates are scheduled more frequently than the specified update frequency. If the 7-Mode system time is ahead of the cluster time, the updates are delayed than the specified update frequency.

Best practices during data copy

To improve SnapMirror replication performance, the source and destination systems should have optimum CPU utilization and memory available. Additionally, network traffic that is unrelated to the migration activity should be minimized so that the throughput is maximized and latency is minimized between the source and

destination systems.

Priority for data transfers

When scheduling the data copy operations, baseline or resynchronization operations take priority over the incremental transfers.

When aborting the data copy operations for releasing the SnapMirror transfers, incremental transfers are aborted first, and then baseline or resynchronization operations are aborted.

For incremental transfers, priority is given to the volumes that are lagging more behind the source volume based on the time elapsed from the previous update.

Editing or deleting a data copy schedule for SnapMirror transfers

You can edit or delete data copy schedules that are used for SnapMirror baseline copy, update, and resynchronization operations, if you have other DR schedules configured or modified that require the transition data copy schedule to be modified.

The project must be in the preparation, data copy, or apply configuration (precutover) phase.

- After a schedule is edited, it takes up to 5 minutes for it to become effective.
- If the throttle limit is changed in the schedule when there are active transfers, the new throttle limit is not applicable for the currently running SnapMirror transfers.

After the current transfer is completed for a particular SnapMirror relationship, a new throttle limit is considered for the subsequent operations for that SnapMirror relationship.

- If you want the throttle limit to be effective immediately for the current SnapMirror transfers, you must pause and resume the project.



Steps



1. From the Dashboard, select a project, and then click Configure Schedule.

All the existing schedules for the project are displayed.

You can also edit or delete schedules from the Configure Schedule option.

2. Edit or delete a schedule:

| If you want to... | Then... |
|---------------------------|--|
| Edit an existing schedule | <div><div>a. Click .</div><div>The Modify Schedule dialog box is displayed.</div><div><div></div><div>You can edit the schedule and SnapMirror parameters for the data copy operations.</div></div><div>b. Make the required changes, and then click Save.</div></div> |

| If you want to... | Then... |
|-------------------|--|
| Delete a schedule | <p>a. Click .</p> <p>The schedule is deleted from the pane.</p> <div>  <p>At least one schedule is required for data transfers. Therefore, you should not delete all of the schedules.</p> </div> |

Example

The following example illustrates how the throttle limit is applied when there are active SnapMirror transfers in the project.

| Schedule | Number of volumes and data copy state | Maximum number of concurrent SnapMirror transfers | Throttle limit | Throttle used by each transfer |
|------------------|--|---|----------------|---|
| Before modifying | Five volumes waiting for baseline | Five | 500 MBps | 100 MBps |
| After modifying | <ul style="list-style-type: none"> For two volumes, baseline is in progress using 100 MBps throttle For three volumes, baseline is completed and are waiting for updates | Five | 250 MBps | <ul style="list-style-type: none"> Two volumes that are in baseline continue to use the 100 MBps throttle Three volumes for which the baseline is completed use the modified 50 MBps throttle for the updates |

When the baseline copy for the two volumes is completed, the new throttle limit of 50 MBps is used for these volume SnapMirror relationships while scheduling the next data copy operations.

Managing logical interfaces

You can add, edit, or remove 7-Mode IP addresses from a project.

The LIFs are configured on the storage virtual machine (SVM) in the apply configuration (precutover) phase:




- Existing 7-Mode IP addresses are created in the administrative “down” state.
- New IP addresses are created in the administrative “up” state.



FC and iSCSI LIFs are not transitioned by the 7-Mode Transition Tool.

Steps

1. Perform one of the following actions:

| If you want to... | Then... |
|-------------------------|--|
| Edit an existing LIF | <ol style="list-style-type: none">a. Click  .b. In the Modify LIF dialog box, make the required changes, and then click Save. |
| Remove an existing LIF | <ol style="list-style-type: none">a. Click  in the table.<div><p>You can remove a LIF only when the project is in the preparation, baseline data copy, pause, precutover, or update state. However, you cannot remove a LIF if the cutover operation has failed.</p></div>b. Click Select 7-Mode LIF.c. Select the 7-Mode IP addresses for transition.d. Enter the target node and target port details.e. Click Save. |
| Add a 7-Mode IP address | <ul style="list-style-type: none">• Click Select 7-Mode LIF.• Select the 7-Mode IP addresses for transition.• Enter the target node and target port details.• Click Save. |
| Add a new LIF | <ol style="list-style-type: none">a. Click Add New LIF below the table.b. Enter the required values.c. Click Save. |

Removing volumes from a project

You can remove the failed volumes from a project during the transition process and abort a volume SnapMirror relationship.

If a volume is removed when the data copy is in progress, then the transfer is not aborted and continues until the operation is completed. The volume is not considered for any further data copy operations.

1. From the Dashboard window, select a project, and then click Volumes.

All of the existing volumes for the project are displayed.

2. Click .

The volume is removed from the project.

The best practice is to delete the ONTAP volumes and release the SnapMirror relationships.

1. From the cluster, delete the ONTAP volumes that are created as part of the transition.

ONTAP 9 commands

2. From the 7-Mode system, release the SnapMirror relationships to the ONTAP volumes.
3. If you abort a volume SnapMirror relationship, perform the following steps:
 - a. Break and delete the SnapMirror relationship between the 7-Mode primary volumes and ONTAP secondary volumes.
 - b. From the primary 7-Mode system, release the SnapMirror relationships to the ONTAP secondary volumes.

Clustered Data ONTAP 8.3 Command Map for 7-Mode Administrators

Pausing and resuming a project

You can pause and resume a project for which data copy has started.

When you pause a primary project, only the copy operation from the 7-Mode primary volume to the corresponding clustered Data ONTAP primary volume is paused. The data copy from the 7-Mode primary volume to the ONTAP secondary volume continues to run according to the schedule. This ensures continued data protection for the primary volumes.

1. From Dashboard, select the project that you want to pause.
2. Click **Pause**.

The data copy operation for all the volumes in the project is stopped.

If the baseline transfer for a volume is in progress and there is no Snapshot copy checkpoint created for the SnapMirror relationships, the pause operation is ignored for the volume. Wait for the Snapshot copy checkpoint to be created and run the pause operation again.

3. Click **Resume**.

The copy operation resumes from the point where it was paused.



The data copy operation is resumed based on the active schedule available at that time.

Related information

Starting baseline data copy

Aborting a project

You might want to abort a running project, for example, when the performance of the controller or the cluster is impacted. You can abort a running project from the Dashboard.

You must consider the following before aborting projects that transition volumes with SnapMirror relationships:

- If both projects have been started, aborting one project aborts the associated project.

For example, if you abort a primary project, the secondary project is also aborted.

- If the abort operation fails on the current project, the associated project is not aborted.
- If only one active project is aborted, the start operation on the associated project fails.
- When you abort a primary project, the copy operation from the 7-Mode primary volume to the ONTAP secondary volume is not aborted.

Only the copy operation from the 7-Mode primary volume to the corresponding ONTAP primary volume is aborted.



If the 7-Mode project is aborted, the only option is to delete the project. It is not possible to resume or restart the project after it is aborted.

Steps

1. From the Dashboard, select the project that you want to abort.
2. Click **Abort**.
3. From the cluster, delete the ONTAP volumes that were created as part of the transition.
4. From the 7-Mode system, release the SnapMirror relationships to the ONTAP volumes.
5. If you abort a volume SnapMirror relationship, you must do the following:
 - a. Break and delete the SnapMirror relationship between the 7-Mode primary volumes and ONTAP secondary volumes.
 - b. From the primary 7-Mode system, release the SnapMirror relationships to the ONTAP secondary volumes.

ONTAP 9 commands

[Data ONTAP 8. 2 Commands: Manual Page Reference for 7-Mode, Volume 1](#)

Deleting a project

If you do not want a project, you can delete it from the Dashboard.

1. From the DashboardProjects Groups pane, click the project group to which the project belongs.
2. From the Group Summary pane, click **Modify**, and then click **Delete**.

You cannot delete a project if the project has failed during the cutover stage.

Troubleshooting issues

You need to be aware of some of the common issues with the 7-Mode Transition Tool and the steps to resolve them.

Downloading transition log files

The 7-Mode Transition Tool creates log files that provide processing details of the transition assessment and migration operations run on your system.

1. Click **Logs** in the top menu.
2. Click **Collect Project Logs** to collect logs related to all of the projects.
3. To collect logs for a given projects, locate the projects from the project list, and then click **Download**.

The logs are downloaded as a .zip file, and the folder name is the timestamp.

Related information

[How to upload a file to NetApp](#)

Log files for the 7-Mode Transition Tool

The 7-Mode Transition Tool creates log files that provide processing details of the transition operations that have occurred on your system. The log files are located in the logs directory of the path where 7-Mode Transition Tool is installed.

You can also use the EMS messages related to SnapMirror logs from the 7-Mode system and the cluster to troubleshoot issues.

The following table lists the log files that are related to a particular transition project:

| Log file path | Contains information about... |
|--|--|
| <i>project_name</i> /transition.log | Debug messages that are specific to a project |
| <i>project_name</i> /zapi-outbound.log | Output of all the Data ONTAP APIs that are executed by 7-Mode Transition Tool for a particular project |

The following table lists the log files that are not related to any particular project:

| Log file path | Contains information about... |
|--------------------|---|
| transition-gui.log | Entries of all the actions performed by using the web interface |

| Log file path | Contains information about... |
|---|---|
| default/audit.log | <ul style="list-style-type: none"> • All the parameters, such as HTTP or HTTPS port and log directory path, that are used by the tool every time 7-Mode Transition Tool is run • All the transition commands that are executed with the outputs |
| default/default/transition.log | Debug messages that are not specific to any project |
| default/STREAM_MANAGEMENT/stream_management.log | Debug messages that are logged by the scheduler while managing the schedules and which do not belong to any project |
| default/default/zapi-outbound.log | Output of all the Data ONTAP APIs that are executed by 7-Mode Transition Tool and which do not belong to any project |
| default/STREAM_MANAGEMENT/zapi-outbound.log | Output of all the Data ONTAP APIs that are executed by the 7-Mode Transition Tool scheduler while managing the schedules and which do not belong to any project |
| server-console.log | Log entries of all the packet exchanges done with the 7-Mode Transition Tool server. This file helps in troubleshooting issues related to a server crash. |

Continuing with the transition if ignorable errors occur

During the transition, you might encounter some errors that block the transition. You can choose to ignore some of these errors by acknowledging the issues through the 7-Mode Transition Tool CLI. You should rerun the failed operation after ignoring the error to continue with the transition.

When you acknowledge an error, it means that you have understood the impact of these errors and acknowledged them.

You must rerun the transition operation after ignoring the error. In some cases, after you acknowledge the issue, Data ONTAP performs corrective actions on the affected aggregates and volumes when the operation is run the next time.

Steps

1. If the transition operation results in any ignorable errors, run the following command from the 7-Mode Transition Tool CLI:

```
transition cbt ignorableerrors add -p project_name -c ignorable_errorcategory
```

ignorable_errorcategory is the type of error that you can ignore.

2. Rerun the transition operation.

The blocking error changes to a warning and the error is shown as `acknowledged`. You can continue the transition with the warning.

Ignorable errors during transition

You might encounter some ignorable errors during the transition. You must acknowledge these errors before continuing with transition.

When you add any ignorable error category to the copy-based transition project by using the 7-Mode Transition Tool CLI, it means that you have understood the impact of the error. You must rerun the transition operation after ignoring the error. At this time, the blocking error changes to a warning message, and the error is shown as “`acknowledged`”. You can continue the transition with the warning.

| Category | Description |
|--|--|
| <code>acknowledge-no-nonascii-filenames-in-source-volumes</code> | Acknowledging this error means that you have verified that the SnapLock volumes that you have selected for the Chain of Custody verification operation do not contain any files that have file names with non-ASCII characters. |
| <code>acknowledge-snaplock-coc-volume-autocommit-period</code> | After this ignorable error is added to the project, the 7-Mode Transition Tool continues with the transition even if the autocommit period property is set on the ONTAP SnapLock volume that is configured to store the results of the SnapLock Chain of Custody verification. |
| <code>nfs-qtrees-exported</code> | Acknowledging this error means that you have understood the differences in the enforcement of the qtree export rules between Data ONTAP operating in 7-Mode and ONTAP. It also means that you have understood the possible manual steps that are required after the NFS exports rules are applied by the 7-Mode Transition Tool. 7MTT Precheck 10111 - How to transition 7-Mode volumes that have qtree level exports |
| <code>ignore-cifs-ad-domain-mismatch</code> | If you acknowledge this error, the 7-Mode Transition Tool continues with the transition of CIFS configurations even if the CIFS Active Directory (AD) domain of the 7-Mode system is different from the CIFS AD domain of the target SVM. You must ensure that the CIFS AD domains of the 7-Mode system and the target SVM are trusted domains. Otherwise, the transition of CIFS configurations to the target SVM fails. How to transition CIFS configurations when Active Directory Domain of CIFS server on 7-Mode and target SVM are different |

| Category | Description |
|--|---|
| <code>ignore-ping-check-from-ic-lifs-to-7mode</code> | After this ignorable error is added to the project, the 7-Mode Transition Tool does not run the check to ping the 7-Mode storage system from the target cluster intercluster LIFs. 7MTT Copy-Based Transition - How to handle the scenario where 7-Mode data-copy IP address is not reachable from target cluster intercluster LIFs |
| <code>ignore-secondary-snapmirror-source-check</code> | After this ignorable error is added to the project, the 7-Mode Transition Tool does not check whether all of the required primary 7-Mode details are added to the secondary project. Due to this error, the 7-Mode Transition Tool might not create the SnapMirror relationships between the primary 7-Mode volumes and secondary cluster volumes during the secondary project complete (cutover) operation. |
| <code>ignore-configuration-limits-check</code> | You can use this option to acknowledge the configuration limits (precheck error) and continue with the transition. By acknowledging this error, you understand that the storage cutover time is expected to take more than 3 to 8 hours, which includes time for the export, halt, and import operations (the automated operations), and also for cabling disk shelves to the new controllers. Storage cutover time considerations for Copy-Free Transition |
| <code>mount-snaplock-volumes-using-7mode-volume-name</code> | After this ignorable error is added to the project, the 7-Mode Transition Tool mounts the SnapLock volumes with the mount path <code>/<7-mode-volume-name></code> |
| <code>acknowledge-no-snaplock-audit-transition-to-mcc</code> | After this ignorable error is added to the project, the 7-Mode Transition Tool continues with the transition even if the target cluster is in a MetroCluster configuration. |

Transitioning a MetroCluster configuration that failed due to switchover or switchback

Transition operations, such as transition prepare, start, pause, resume, or complete, fail during a MetroCluster switchover or after a MetroCluster switchback. You must then manually reestablish the SnapMirror relationships to resume transition.

Steps

1. Wait for the switchback to complete.
2. From the cluster CLI, use the `snapmirror create` command to reestablish the SnapMirror relationships between the 7-Mode volumes and clustered Data ONTAP volumes that were part of the transition.

3. From the 7-Mode Transition Tool, rerun the transition operation.

Data copy operations start when the next schedule becomes active.

Cannot select a secondary volume from the Volume selection pane

You cannot select a secondary volume from the Volume selection pane and transition it as a volume SnapMirror relationship if the volume is not online or if the SnapMirror relationship is broken.

Workaround

Use the ONTAP command-line to perform either of the following workarounds:

- Bring the secondary volume online if it is offline.
- Fix the SnapMirror relationship if it is broken.

Cannot select a volume for transition if the tool fails to retrieve the volume information

The Volume selection pane displays the `Failed to retrieve volume information` error message when you try to select a volume. This error usually occurs if the controller is busy.

Workaround

Wait until the controller is not busy, and then select the volume again.

Cannot proceed from the Volume mapping pane if the selected SVM has no aggregate

The tool displays an error message, `SVM does not have aggregate assigned. Please assign some aggregates to the SVM and click on the refresh button`, on the **Volume mapping** pane if the selected SVM has no aggregate other than the root node aggregate.

Workaround

Add an aggregate to the SVM in clustered Data ONTAP.

Compression is not enabled after transition from Data ONTAP 7.3.x

If you transition a 32-bit volume on a system running Data ONTAP 7.3.x to an ONTAP volume that is in a 64-bit aggregate, compression is not enabled on the ONTAP volume after the transition. Compression cannot be enabled until the volume is upgraded from 32-bit to 64-bit.

Workaround

1. Wait for the upgrade to finish.

You can use the `volume show -fields block-type -volume vol_name` command to verify that the block type of the volume changes to 64-bit.

2. Enable compression by using the `volume efficiency modify` command.

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.