




7-Mode Transition Tool 2.1

Data and Configuration Transition Guide

For Transitioning to Clustered Data ONTAP®



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-10060_A0
June 2015

Contents

Transition overview	7
What the 7-Mode Transition Tool is	7
Features and volumes not supported for transition	8
Unsupported 7-Mode SAN features	10
Transition terminology	11
Limits for transition	12
Installing, upgrading, or uninstalling the 7-Mode Transition Tool	13
System requirements for 7-Mode Transition Tool	13
Installing the 7-Mode Transition Tool on Windows	15
Upgrading the 7-Mode Transition Tool on Windows	15
Uninstalling the 7-Mode Transition Tool on Windows	16
Enabling login for users not part of the Administrator group	16
Logging in to the tool	17
Installing 7-Mode Transition Tool on Linux	18
Uninstalling 7-Mode Transition Tool on Linux	19
Commands for managing 7-Mode Transition Tool service on Linux systems	20
Modifying the configuration options of the 7-Mode Transition Tool	20
Collecting and assessing the inventory information	22
Storage and host version requirements for transition assessment	22
Preparing the 7-Mode systems and hosts for transition assessment	22
Assessing controllers and hosts	23
Generating an assessment report by adding systems to the 7-Mode Transition Tool	24
Generating an assessment report by importing the inventory report XML ...	26
How you can use the assessment executive summary for transition assessment	26
Migration process and phases	29
How you transition a stand-alone volume	30
How you transition volumes in a SnapMirror relationship	34
NFS migration workflow	39
CIFS migration workflow	40
FC migration workflow	40
iSCSI migration workflow	41
Preparing for migration	43
Data ONTAP version requirements for migration	43
License requirements for transition	43
7-Mode Transition Tool service must be running	44
Restrictions for transition	44
Preparing the cluster for transition	45
Preparing the 7-Mode system for transition	47
Port requirements for communicating with the 7-Mode Transition Tool	49
Prerequisites for transitioning NFS configurations	49

Prerequisites for transitioning CIFS configurations	50
Prerequisites and considerations for transitioning 7-Mode volumes to a MetroCluster configuration in clustered Data ONTAP	51
Preparation for SAN transition	52
Preparing for transition to clustered Data ONTAP 8.3 and later	52
Expanding an aggregate to the 64-bit format	54
Finding and removing 32-bit data from source volumes and Snapshot copies	55
Migrating data and configuration from 7-Mode volumes	57
Transition preparation checklist	58
Adding controllers and clusters	62
Creating a transition project	63
Running prechecks	66
Severity levels for precheck messages	67
Starting baseline data copy	68
Applying 7-Mode configurations	69
Applying 7-Mode configurations for volumes without any LUNs (NAS only)	69
Applying configurations on volumes that have at least one LUN	71
Completing a transition project	72
Managing a transition project	74
Editing a subproject	74
Managing SnapMirror transfers and schedule	74
Creating a data copy schedule for SnapMirror transfers	74
Considerations for creating a data copy schedule	79
Editing or deleting a data copy schedule for SnapMirror transfers	80
Performing on-demand SnapMirror updates	81
Transitioning volumes by excluding local user and groups or netgroup configurations	82
Managing logical interfaces	83
Removing volumes from a subproject	83
Pausing and resuming a subproject	84
Aborting a subproject	85
Deleting a project	86
Transitioning volumes using the 7-Mode Transition Tool CLI	87
Scenarios for selecting volumes in a session	87
Transitioning volumes by using the 7-Mode Transition Tool command-line interface	88
Creating a transition session	88
Running prechecks	91
Starting data copy for a transition session	92
Applying configurations to the clustered Data ONTAP volumes	93
Completing the transition	94
Commands for managing transition	94
Considerations for transition	98

Considerations for using SnapMirror for transition	98
Guidelines for deciding when to perform cutover	99
NFS only configuration transition	99
How NFS exports are transitioned	100
Considerations for transitioning UNIX users and groups	102
Considerations for transitioning CIFS local users and groups	103
Space considerations when transitioning SAN volumes	104
Considerations for deduplication and compression	104
Considerations for FlexClone volumes	105
Considerations for transitioning 7-Mode volumes to Storage Virtual Machines (SVMs) in the nondefault IPspace	105
Considerations for Snapshot copies	105
Considerations for quotas	105
Considerations for transitioning LDAP and DNS configurations	106
Impact of takeover and giveback on transition	106
About backing up volumes in transition to tape	106
Performing manual post-transition tasks	108
Recovering from a failed LUN transition	108
Viewing transitioned SAN configurations	109
Limitations with 7-Mode Snapshot copies of LUNs managed by SnapDrive and SnapManager after transition	110
Configuring cron job schedules on the remote site after transitioning a MetroCluster configuration	111
Deleting 7-Mode volume Snapshot copies from transitioned volumes	111
Consolidating Snapshot policies and cron schedules from transitioned volumes ...	112
Troubleshooting issues	113
Collecting tool logs	113
Log files for the 7-Mode Transition Tool	113
Tool installation fails with an incompatible OS error	114
Installation or uninstallation of the 7-Mode Transition Tool is blocked	114
Reinstalling or upgrading the 7-Mode Transition Tool fails on Windows system ..	115
Windows Program Compatibility Assistant notification displayed while installing the tool	115
Incorrect Data ONTAP version displayed if the 7-Mode system is upgraded after project creation	115
Transitioning a MetroCluster configuration that failed due to switchover or switchback	116
Cannot select a secondary volume from the Volume selection pane	116
Cannot select a volume for transition if the tool fails to retrieve the volume information	116
Cannot proceed from the Volume mapping pane if the selected SVM has no aggregate	117
Cannot use 7-Mode Transition Tool for transitioning 7-Mode volumes to Storage Virtual Machines (SVMs) in nondefault IPspace	117
Copyright information	118

Trademark information 119

**How to send comments about documentation and receive update
notifications 120**

Index 121

Transition overview

Transitioning to clustered Data 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, as well as making necessary environmental updates.

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 clustered Data 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 clustered Data ONTAP. For example, a 7-Mode volume should be mapped to a Storage Virtual Machine (SVM) and a volume in clustered Data 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 clustered Data ONTAP system.

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

Related information

[NetApp Documentation: Data ONTAP 8 \(current releases\)](#)

[NetApp Technical Report 3982: NetApp Clustered Data ONTAP 8.2 Introduction](#)

What the 7-Mode Transition Tool is

The 7-Mode Transition Tool enables you to assess 7-Mode controllers, hosts, and applications for transition and perform copy-based transitions of FlexVol volumes and configurations from Data ONTAP operating in 7-Mode to new hardware that is running clustered Data ONTAP 8.2.x and 8.3.x, with minimum client disruption and retention of storage efficiency options.

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. It is a best practice to enable TLS and disable SSLv3 on the storage system in order to avoid SSLv3 security vulnerabilities.

The 7-Mode Transition Tool works at the volume level. It uses SnapMirror technology for transferring data by setting up a SnapMirror relationship known as *transition data protection (TDP)* between the 7-Mode volumes and the clustered Data ONTAP volumes. SnapMirror relationships between 7-Mode and clustered Data ONTAP are one-way and are supported only for the purpose and duration of a migration that is part of a transition to clustered Data ONTAP.

Note: It is not possible to resynchronize data from a clustered Data ONTAP volume back to a 7-Mode volume. You can contact your NetApp representative for host-based tools that can copy data from clustered Data ONTAP to Data ONTAP operating in 7-Mode.

The 7-Mode Transition Tool enables you to perform the following tasks:

- Collect information about 7-Mode controllers, hosts, and applications and assess the current features and functionality to identify compatibility concerns with the clustered Data ONTAP version selected for transition.

- Transition NAS and SAN data and configurations to clustered Data ONTAP.
Note: SAN transition is supported only to clustered Data ONTAP 8.3 and later.
- Migrate data and configuration from systems running different 7-Mode versions to clustered Data ONTAP versions 8.2.x or 8.3.x.
- Transition existing stand-alone FlexVol volumes from a 7-Mode controller to a Storage Virtual Machine (SVM, formerly known as Vserver).
- Transition both the primary and the secondary FlexVol volumes in a SnapMirror relationship, and retain the SnapMirror relationship between the transitioned volumes that are on a cluster.
- Preserve Snapshot copies, deduplication, and compression savings in the volumes after the transition.
- Use the 7-Mode deployments while data is being copied to a newly deployed cluster.
- Apply configurations for NAS volumes in clustered Data ONTAP before disconnecting client access.
 This enables you to verify the configurations in clustered Data ONTAP and to minimize downtime.

The tool provides both a web interface and a command-line interface.

Related concepts

[Data ONTAP version requirements for migration](#) on page 43

Features and volumes not supported for transition

You cannot transition certain 7-Mode volumes, such as traditional volumes or SnapLock volumes, by using the 7-Mode Transition Tool. Also, you cannot transition certain 7-Mode features because either some features might not be available in clustered Data ONTAP or the tool cannot transition these features.

Transitioning a 7-Mode volume to an SVM with Infinite Volume is not supported. You can transition a 7-Mode volume only to an SVM with FlexVol volumes.

You cannot transition the following 7-Mode volumes:

- Restricted or offline volumes
- Traditional volumes
- Volumes with NFS-to-CIFS character mapping (charmap)
 If the target cluster is running Data ONTAP 8.3 or later, you can transition volumes with charmap enabled.
- Volumes with Storage-Level Access Guard configurations
 If the target cluster is running Data ONTAP 8.3 or later, you can transition volumes with this configuration.
- Volumes that contain qtrees with Storage-Level Access Guard configurations
 If the target cluster is running Data ONTAP 8.3.1 or later, you can transition volumes that contain qtrees with this configuration.
- SnapLock volumes
- FlexCache volumes

- 32-bit volumes and 64-bit volumes that have 32-bit Snapshot copies if the destination cluster is running Data ONTAP 8.3 or later
- FlexClone volumes
FlexClone volumes can be transitioned as FlexVol volumes, but the clone hierarchy and storage efficiency will be lost.
- Root volume of a vFiler unit, where the root volume is based on a qtree that belongs to the default vFiler unit

You cannot transition the following 7-Mode configurations:

- Export rules with Kerberos security:
 - krb5p if the target cluster is running Data ONTAP 8.3 and later
 - krb5i and krb5p if the target cluster is running Data ONTAP 8.2.x
- IPv6 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 relationship is broken or deleted.
[*NetApp Technical Report 4052: Successfully Transitioning to Clustered Data ONTAP \(Data ONTAP 8.2.x and 8.3\)*](#)
- FPolicy configurations
- BranchCache configurations
- Antivirus configurations
- PC-NFS
- WebNFS
- NFSv2
- FTP
- Sub-volume NFS exports other than qtree-level NFS exports
Note: Transition of qtree-level NFS exports is supported if the target cluster is running Data ONTAP 8.2.1 and later.
- Fencing of NFS clients from one or more file system paths
- Transition of volumes serving CIFS data that is being accessed by local users and groups to a Data ONTAP 8.2 destination cluster
7-Mode Transition Tool supports the transition of local users and groups only to clustered Data ONTAP 8.2.1 and later.
- CIFS NT4 authentication
- CIFS NetBIOS aliases
- CIFS share-level ACLs with UNIX-style permissions
- Configuration to allow the administrator to connect to the CIFS home directories of other users (cifs.homedirs_public_for_admin option on the 7-Mode system)
- Synchronous SnapMirror configurations

- Qtree SnapMirror relationships

Although qtree SnapMirror source volumes can be migrated, you must first break and delete the qtree SnapMirror relationship. If space is available on a 7-Mode storage system (preferably the source system), the staging method can be used to break out qtrees into independent volumes. After being staged into volumes in a 7-Mode system, you can use the 7-Mode Transition Tool to migrate these volumes directly to clustered Data ONTAP.

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

- Disaster recovery vFiler unit (vFiler DR)

However, to manually transition the volumes that are part of a vFiler DR relationship, you must first delete the disaster recovery capability from the destination vFiler unit and then transition the destination and source volumes as a volume SnapMirror relationship.

- NDMP configurations
- SFTP server
- TFTP server
- GARP VLAN Registration Protocol (GVRP)
- NIS for host name lookup
- Settings in the `/etc/hosts.equiv` file

You can work with your NetApp representative to identify workarounds for transitioning some of these features.

Related concepts

[*Restrictions for transition*](#) on page 44

Unsupported 7-Mode SAN features

You should be aware of the 7-Mode SAN features that are not supported in clustered Data ONTAP so that you can take any necessary actions before the transition.

Transition fails with the following 7-Mode SAN features:

- Snapshot copy-backed LUN clones
You must split or delete your Snapshot copy-backed LUN clones prior to transition.
- LUNs with `ostype` of **vld** and **image**
You must delete LUNs with `ostype` of **vld** or **image** prior to transition.
- LUN clone split
Active LUN clone splits must finish or be aborted.

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

- The `lun share` command
Sharing a LUN over NAS protocols is not supported in clustered Data ONTAP.
- SnapValidator

The following 7-Mode feature is not transitioned by the 7-Mode Transition Tool:

- An igroup bound to a portset
Portsets are not transitioned. You must manually configure igroups that are bound to a portset after 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 clustered Data 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 clustered Data ONTAP version selected for transition.

Migrate

Refers to transitioning data and configurations from the 7-Mode volumes to clustered Data 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.

Subproject

In the 7-Mode Transition Tool, a subproject is a container within a project that groups a set of volumes for transition. A subproject also contains the configuration details that are required for the transition, such as the SVM mapping, volume mapping, interface mapping, and data copy schedules.

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

The term *session* that is used in the 7-Mode Transition Tool CLI means the same thing as *subproject*.

Stand-alone subproject

Volumes in this subproject 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 subproject

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

Secondary subproject

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

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 clustered Data 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 a clustered Data 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 a clustered Data 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 subproject and number of subprojects 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	4
Number of hosts	20

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

Parameter	Maximum limit
Number of volumes in a subproject	160
Recommended limit for total number of volumes across all active subprojects in the web interface at any given time	240
Maximum number of active projects at any given time	50 (each project can have either a stand-alone subproject or a primary and a secondary subproject) Note: If you have fifty active projects, you can replicate data from fifty 7-Mode storage systems and their associated SnapMirror destinations.
Number of data copy schedules per subproject	7

Installing, upgrading, or uninstalling the 7-Mode Transition Tool

You can download and install the 7-Mode Transition Tool on a Windows or Linux system. On a Windows system, you can also upgrade to 7-Mode Transition Tool 2.1 from version 2.0. When required, you can uninstall the 7-Mode Transition Tool.

On Windows, you can install 7-Mode Transition Tool by using the standard wizard-based installer. On Linux, you must install 7-Mode Transition Tool by using the command-line interface.

System requirements for 7-Mode Transition Tool

You must ensure that you have the required configuration to install and run the 7-Mode Transition Tool.

Hardware requirements

- Dual-core x64 processor, 1.0 GHz or more
- 4-GB RAM
- 40-GB HDD

Windows system: software requirements

- Your 64-bit Windows system must be running one of the following:
 - Windows 7 Enterprise
 - Windows 7 Enterprise SP1
 - Windows Server 2008 Standard
 - Windows Server 2008 R2 Standard
 - Windows Server 2008 R2 Enterprise with SP1
 - Windows Server 2012 Standard
 - Windows Server 2012 R2 Standard
 - Windows 8.1 Enterprise
- 64-bit Oracle Java Runtime Environment (JRE) 1.8 update 45

Note: If the Windows host does not have JRE or has 32-bit JRE installed, the 7-Mode Transition Tool installer automatically installs 64-bit JRE 1.8 update 45. If an earlier version of 64-bit JRE is installed, the installer automatically updates it to JRE 1.8 update 45.

Linux system: software requirements

The 7-Mode Transition Tool web interface is not supported on Linux. You must use the CLI to perform the transition tasks on a Linux system.

- Your Linux system must be running one of the following:
 - Red Hat Enterprise Linux 5.6

- Red Hat Enterprise Linux 6.0
- CentOS release 6.4
- 64-bit Oracle JRE 1.8 update 45
- 7-zip utility (rpmforge-release-0.3.6-1.el5.rf.x86_64.rpm)

Note: JRE 1.8 and the 7-zip utility must be installed on the Linux system before configuring 7-Mode Transition Tool.

Server configuration requirements

To access the web interface, you must ensure that the Windows system on which the tool is installed is configured as follows:

- Ports 8443 (for HTTPS) and 8088 (for HTTP) must be available.
If port 8443 is not available, or if you want to use a different port for HTTPS, you must change the port specified by using the `tool.https.port` parameter in the `transition-tool.conf` file. For an alternative HTTP port, you must change the port specified by the `tool.http.port` parameter in the `transition-tool.conf` file.
Note: You must restart the 7-Mode Transition Tool service after changing the port in the configuration file.
- The firewall should be either turned off or configured to allow traffic on the port that is used to access the tool.

Client (web interface) requirements

The system that accesses the web interface must have the following:

- One of the following web browsers:
 - Microsoft Internet Explorer 9, 10, or 11
 - Google Chrome 27 or later
 - Mozilla Firefox 20 or later

All browsers must be TLS enabled.

Note: If you are using Microsoft Internet Explorer as your web browser, you must disable SSLv2.

- A screen resolution of 1280 × 1024 or higher

For more information, see the Interoperability Matrix.

Note: Every time you install a new version of the tool, you must clear the browser cache by pressing Ctrl+F5 on the system.

Related concepts

Port requirements for communicating with the 7-Mode Transition Tool on page 49

Related information

Interoperability Matrix Tool: mysupport.netapp.com/matrix

Installing the 7-Mode Transition Tool on Windows

You can install the 7-Mode Transition Tool on your Windows system by using the wizard-based installer.

Before you begin

- You must have reviewed the [System requirements for 7-Mode Transition Tool](#).
- You must have the necessary administrator privileges to install and launch the application.
- You must have downloaded the .exe file from the NetApp Support Site.
- Every time you install a new version of the tool, you must have cleared the browser cache by pressing Ctrl+F5.
- You must complete any in-progress subprojects before you uninstall the earlier version of the tool.

Steps

1. Run the NetApp_7ModeTransitionTool_Setup.exe file.
2. In the **7-Mode Transition Tool Setup** window, follow the on-screen prompts.
3. Optional: Select **Launch the tool** to open the 7-Mode Transition Tool immediately after the installation.
4. Click **Finish** to complete the installation.

Related concepts

[Data ONTAP version requirements for migration](#) on page 43

[License requirements for transition](#) on page 43

[Port requirements for communicating with the 7-Mode Transition Tool](#) on page 49

Related information

[NetApp Support Site: mysupport.netapp.com](#)

Upgrading the 7-Mode Transition Tool on Windows

You can upgrade to this release of the 7-Mode Transition Tool on your Windows system without uninstalling the previous version. Upgrading enables you to view the transition history of completed and aborted projects that were present in previous version of the tool.

Before you begin

- You must have reviewed the [System requirements for 7-Mode Transition Tool](#).
- You must have the necessary administrator privileges to install and launch the application.
- All projects must be in completed or aborted state.
 - Note:** After the upgrade, all pre-existing projects can only be viewed; they cannot be restarted.
- All instances of the tool must be closed.

About this task

If you are using 7-Mode Transition Tool 2.0, you can upgrade to 7-Mode Transition Tool 2.1.

Steps

1. Run the `NetApp_7ModeTransitionTool_Setup.exe` file.
2. In the **7-Mode Transition Tool Setup** window, follow the on-screen prompts.
The required action is automatically selected as **Upgrade**.
3. Optional: Select **Launch the tool** to open the 7-Mode Transition Tool immediately after the upgrade.
4. Click **Finish**.

Related information

[NetApp Support](#)

Uninstalling the 7-Mode Transition Tool on Windows

You can uninstall the 7-Mode Transition Tool by using the Windows Uninstall program utility.

Before you begin

- You must have completed all in-progress transitions.
Caution: Installing a newer version or reinstalling the same version of the tool does not allow you to access the transition operations started by the uninstalled instance.
- If the 7-Mode Transition Tool is running, it must be closed.

About this task

JRE is not removed as part of the uninstallation.

Steps

1. Click **Start > Control Panel > Uninstall a program**.
2. Select the 7-Mode Transition Tool from the list of programs and click **Uninstall**.
3. Complete the wizard to uninstall the 7-Mode Transition Tool.
You can choose to save the subproject details and logs for troubleshooting purposes.

After you finish

You must make a backup of the log files if you choose to retain the log files during uninstallation. The 7-Mode Transition Tool installer installs the tool to the same folder and overwrites the log files the next time you run it.

Enabling login for users not part of the Administrator group

By default, users must be members of the Administrator group on the Windows system on which the 7-Mode Transition Tool is installed. At any time, you can enable login for users who are not

members of the Administrators group and who cannot install the tool but are authenticated to use the Windows system on which the tool is installed.

About this task

For users who are not members of the Administrator group, login access and a valid user name must be provided by a user who has admin privileges (is a member of the Administrator group). The newly added user is provided access to all the projects in the tool.

Steps

1. Enable users who are not part of the Administrator group to log in to the tool by performing the following steps:
 - a. Log in to the Windows host by using admin privileges.
 - b. Set the value of the `tool.login.non.admin.enabled` parameter in the `transition-tool.conf` file to **true**.
By default, the value is set to **false**.
 - c. Save the file.
 - d. Restart the transition server:
service transition-service restart
2. Create a valid user name for the user who wants to log in to the tool but is not a member of the Administrator group.

Logging in to the tool

To transition volumes using the 7-Mode Transition Tool (web interface), you must log in to the tool using your Windows credentials. The procedure to log in varies depending on the type of user credentials you have.

Before you begin

If you are a user who is not a member of the Administrator group, you must be provided access to log in to the tool by a member of the Administrator group. You must also be provided with a valid user name by a member of the Administrator group.

About this task

Guest users are not allowed to log in to the tool.

Steps

1. Navigate to the login page:

If you are a...	Then...
User who is a member of the Administrator group and if you are accessing the tool (web interface) from the server on which it was installed	Double-click the 7-Mode Transition Tool icon on your desktop.

If you are a...	Then...
User who is a member of the Administrator group and if you are accessing the tool (web interface) from another computer	Use a supported browser to navigate to one of the following URLs: • <code>https://IP_address:port/transition</code> • <code>http://IP_address:port/transition</code> <i>IP_address</i> is the IP address of the server and <i>port</i> can be 8443 (default), the value specified for <i>tool.https.port</i>, or 8088, the value specified for the <i>tool.http.port</i> parameter in the <i>transition-tool.conf</i> file.
User who is not a member of the Administrator group and irrespective of whether you are accessing the tool (web interface) from the server on which it was installed or from another computer	Use a supported browser to navigate to one of the following URLs: • <code>https://IP_address:port/transition</code> • <code>http://IP_address:port/transition</code> <i>IP_address</i> is the IP address of the server and <i>port</i> can be 8443 (default), the value specified for the <i>tool.https.port</i>, or 8088, the value specified for the <i>tool.http.port</i> parameter in the <i>transition-tool.conf</i> file.

Members of Administrator group who have not installed the tool but want to use it have to access the login page by navigating to `https://IP_address:port/transition`.

2. If the browser displays a message related to the self-signed SSL certificate, accept the self-signed certificate to continue.

If you are using the 8088 port (`http://IP_address:port/transition`), then the self-signed SSL certificate message is not displayed.

3. Log in to the application by using your Windows credentials.

Installing 7-Mode Transition Tool on Linux

You can install the 7-Mode Transition Tool on your Linux system by using the command-line interface.

Before you begin

- You must have reviewed the [System requirements for 7-Mode Transition Tool](#).
- You must have the root user privileges to install and launch the application.
- You must have downloaded the `tar .gz` file from the NetApp Support Site.
- 64-bit JRE 1.8 update 45 must be installed on the Linux system.
- 7-zip utility (`rpmforge-release-0.3.6-1.el5.rf.x86_64.rpm`) must be installed.

Steps

1. Extract the `tar .gz` file to a directory in the Linux system:


```
tar -xzf NetApp_7ModeTransitionTool_Setup.tar.gz -C directory_path
```
2. Change the directory to the installation directory:


```
cd directory_path/NetApp_7ModeTransitionTool
```
3. Configure the 7-Mode Transition Tool server:

./configure

This starts the 7-Mode Transition Tool service on the Linux system.

Example

```
bash-4.2# tar -xzvf NetApp_7ModeTransitionTool_Setup.tar.gz -C /  
root/Downloads/extracted_folder  
bash-4.2# cd /root/Downloads/extracted_folder/  
NetApp_7ModeTransitionTool  
bash-4.2# ./configure
```

Related information

[NetApp Support](#)

Uninstalling 7-Mode Transition Tool on Linux

You can uninstall the 7-Mode Transition Tool on Linux by using the command-line interface.

Before you begin

- You must have completed all in-progress transitions.
Caution: Installing a newer version or reinstalling the same version of the tool does not allow you to access the transition operations started by the uninstalled instance.
- If the 7-Mode Transition Tool is running, it must be stopped.

About this task

JRE is not removed as part of the uninstallation.

Steps

1. To uninstall 7-Mode Transition Tool, run the following command from the location where you have extracted the 7-Mode Transition Tool:

./unconfigure

2. Remove the 7-Mode Transition Tool directory by using the following commands:

a. **cd ..**

b. **rm -rf NetApp_7ModeTransitionTool**

Example

```
bash-4.2# cd /root/Downloads/extracted_folder/  
NetApp_7ModeTransitionTool  
bash-4.2# ./unconfigure  
bash-4.2# cd ..  
bash-4.2# rm -rf NetApp_7ModeTransitionTool
```

Commands for managing 7-Mode Transition Tool service on Linux systems

You can configure, manage, and unconfigure the 7-Mode Transition Tool service on Linux systems by using certain Linux commands.

The 7-Mode Transition Tool service is registered as transition-service on the Linux system.

If you want to...	Use this command...
Configure the 7-Mode Transition Tool service	<code>./configure</code>
View the status of the transition service	<code>service transition-service status</code>
Stop the transition service	<code>service transition-service stop</code>
Start the transition service	<code>service transition-service start</code>
Restart the transition service	<code>service transition-service restart</code>
Unconfigure the 7-Mode Transition Tool service	<code>./unconfigure</code>

Modifying the configuration options of the 7-Mode Transition Tool

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. This file contains information about all the configurable options that are necessary for the tool to operate.

About this task

The file contains various options; for example, you can specify the port on which the tool service starts, and the port that the tool uses to communicate with the 7-Mode system or cluster.

Steps

1. From the system where the 7-Mode Transition Tool is installed, open the `\etc\conf\transition-tool.conf` file from the install directory and edit it.

Example

Port 8443 or 8088 must be available for the 7-Mode Transition Tool to communicate with the 7-Mode system and cluster. If port 8443 or 8088 is not available or if you want to use a different port, you must change the port specified by the `tool.https.port` option in the `transition-tool.conf` file. If port 8088 is not available or if you want to use a different port, you must change the port specified by the `tool.http.port` option in the `transition-tool.conf` file.

2. Restart the 7-Mode Transition Tool service for the new values to take effect.

If your system type is...	Do the following...
Windows	a. Click Start > Control Panel > System and Services > Administrative Tools > Services. b. Right-click the Transition Tool service and click Stop. c. Right-click the Transition Tool service again and click Start.

If your system type is...	Do the following...
Linux	Run the following command: service transition-service restart

3. After the 7-Mode Transition Tool service is restarted, you must add the 7-Mode system and cluster credentials again.

Collecting and assessing the inventory information

You can collect inventory information from controllers and hosts, assess the features and functionalities of these systems, and identify how these features and functionalities work in the clustered Data ONTAP version selected for transition.

You can either use 7-Mode Transition Tool to collect the inventory information or import the inventory XML report generated by Inventory Collect Tool and then perform the assessment.

Storage and host version requirements for transition assessment

You must be aware of the versions of Data ONTAP operating in 7-Mode and the hosts that are supported for transition assessment.

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

Related information

[Interoperability Matrix Tool: mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix)

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
Important: 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

Important: 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 by entering the following command:

```
secureadmin setup -f ssh
```

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

- b. Enable SSH by entering the following command:

```
secureadmin enable ssh2
```

- c. Enable password authentication on the SSH server by entering the following command:

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

- d. Enable SSH access to the host by entering the following command:

```
options ssh.access
```

For more information about all the 7-Mode commands, see the man pages.

5. Prepare your Windows host systems for the following:

- Enable WMI access.
For more information about enabling WMI access, see the host documentation.
- 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, then the user name 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 are reachable from the Windows system from which the 7-Mode Transition Tool is run.

Related information

[Documentation on the NetApp Support Site: mysupport.netapp.com](http://mysupport.netapp.com)

Assessing controllers and hosts

You can use 7-Mode Transition tool to assess host and controllers by using previously collected inventory information about the systems generated by the Inventory Collect Tool. Alternatively, you

can use the 7-Mode Transition tool to collect inventory information by adding systems to the tool and then perform the assessment.

About this task

- You should not perform both assessment and migration operations on a controller simultaneously.
- You should avoid performing assessment operations on active storage controllers during peak hours.

Choices

- [Generating an assessment report by adding systems to the 7-Mode Transition Tool](#) on page 24
- [Generating an assessment report by importing the inventory report XML](#) on page 26

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

You can collect inventory information for controllers and hosts 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 clustered Data ONTAP version selected for transition.

Before you begin

- The user name for the storage system and hosts must have sufficient privileges to execute the commands listed in the `readme` file.
The file is located at `7-Mode Transition Tool_installed_location\bin\ict`.
- You must have prepared the 7-Mode systems and hosts for transition assessment.
- For assessing Windows system, you must have a domain user account.
- If you are adding multiple systems for assessment, you must have created a text file that is encoded in ASCII format and contains the system details, one system per line.
Each system details must be in the following format:

```
(ontap|windows|vmware|linux)://[(<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.
- All host names and storage system configurations, such as CIFS share names, user names, and group names, must be in the UTF-8 format.

About this task

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 systems must be added to the tool again.

Steps

1. Log in to the 7-Mode Transition Tool, and in the **Collect and Assess** section, click **Get Started**.
2. Click **Add Systems**, and then perform one of the following actions in the **Add System** window:

If you want to add the credentials for...	Then...
A single system	a. Enter the name or the IP address of the system. b. Enter the user name and password for the specified system. The user name for the storage system must have sufficient administrative privileges to collect inventory information. c. Select the system type. For example, if you want to add a 7-Mode storage system, you must set the type as Data ONTAP storage.
Multiple systems	a. Click Browse, and then select 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. Assess the controllers and hosts by generating the assessment report:

- a. Select the systems for which you want to perform 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. Optional: 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.

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

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 **n/a**.

5. Optional: View the processing details of the transition assessment operation by clicking **Migrate > Logs > Collect Tool Logs**.

Related tasks

[Collecting tool logs](#) on page 113

Generating an assessment report by importing the inventory report XML

You can import the inventory XML report generated by the Inventory Collect Tool and then assess features and functionalities of hosts and controllers and identify how these work in the clustered Data ONTAP version selected for transition by creating an assessment report.

Before you begin

- You must have run the Inventory Collect Tool and generated the inventory report XML file.
- You must have prepared the 7-Mode systems and hosts for transition assessment.

About this task

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 7-Mode Transition Tool, and in the Collect and Assess section, click **Get Started**.
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. Optional: 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 in Microsoft Excel and AssessmentExecutiveSummary in Microsoft Word by using Microsoft Office 2007 or later versions.
To view the AssessmentWorkbook in Microsoft Excel, you might have to enable macros in Excel.
11. Optional: View the processing details of the transition assessment operation by clicking **Migrate > Logs > Collect Tool Logs**.

Related tasks

[Collecting tool logs](#) on page 113

How you can use the assessment executive summary for transition assessment

The transition executive summary provides a summary of the 7-Mode controllers and hosts 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 clustered Data ONTAP version of the target cluster that you selected during assessment.

Data collection summary

You can view the list of 7-Mode controllers and hosts for which you have collected information. You can view the Data 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 Transition 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 SnapLock volumes, traditional volumes, or FlexCache volumes because these features are not supported in clustered Data ONTAP.

For most of the configurations, 7-Mode Transition Tool is the recommended tool for transition. However, there are some workloads that cannot be transitioned by 7-Mode Transition Tool for which 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, they have different NIS configurations, they have the same volume names or CIFS share names, and how many local users and groups are there 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.

Migration process and phases

The data and configuration migration process using the 7-Mode Transition Tool consists of the following phases: preparation, baseline data copy, precutover, and storage cutover. You should understand the phases to manage the transition effectively.

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 clustered Data ONTAP. It checks that the cluster is configured properly and is able to support the transition. Any errors must be resolved before continuing with the transition. Although the tool allows you to continue without resolving warnings, it is a best practice to address any warnings before continuing with the migration. You can run the prechecks multiple times to verify that all 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 (clustered Data ONTAP) systems. The assessment step only evaluates the migration source systems, checking feature and functionality differences with clustered Data ONTAP.

Baseline data copy

New volumes are created on the SVM, a SnapMirror relationship is established between the 7-Mode and clustered Data ONTAP volumes, and a baseline transfer is performed. After the baseline is complete, incremental transfers are automatically run according to a user-defined 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%).

Precutover

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

Tip: Although incremental updates are not required, it is a best practice to perform an incremental transfer as close to the storage cutover as possible so as 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 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, such as applying SAN configurations and quotas, that were not transitioned during precutover are also applied to the SVM at this time.

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

How you transition a stand-alone volume

Transitioning a stand-alone volume includes different phases: preparation, data copy, 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.



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. Collecting and adding 7-Mode storage system and volume information
2. Running the transition precheck
3. Collecting and adding cluster, SVM, and aggregate information
4. Collecting 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

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

5. Creating data copy schedules for baseline copy and incremental updates

You should 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 clustered Data ONTAP volumes. The 7-Mode Transition Tool performs the following tasks in this phase:

1. Creating the clustered Data ONTAP with read-only access
2. Setting up a transition peer relationship between the 7-Mode system and the SVM
3. Establishing a transition SnapMirror relationship (relationship of type TDP) between the 7-Mode volumes and clustered Data ONTAP volumes

4. Completing the baseline data copy transfer based on schedule inputs
5. Performing scheduled incremental updates to the clustered Data ONTAP volumes

Precutover phase

In this phase, configurations from the 7-Mode volumes are copied to the clustered Data ONTAP volumes. There are two modes for the precutover phase: precutover read-only and precutover read/write. You should create separate subprojects for NAS and SAN volumes because precutover read/write mode is disabled if you have at least one SAN volume in your subproject.

Subproject with volumes that do not have LUNs (NAS only)

If a subproject does not have any volumes with LUNs, both the precutover read-only and precutover read/write modes are supported.

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

1. Performing an incremental update from 7-Mode volumes to clustered Data ONTAP volumes
2. Breaking the SnapMirror relationship between 7-Mode volumes and clustered Data ONTAP volumes
3. Collecting configurations from 7-Mode volumes, and applying the configurations to the clustered Data ONTAP volumes and SVM
4. Configuring 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. Resynchronizing the SnapMirror relationship between 7-Mode volumes and clustered Data ONTAP volumes

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

1. Performing an incremental update from 7-Mode volumes to clustered Data ONTAP volumes
2. Breaking the SnapMirror relationship between 7-Mode volumes and clustered Data ONTAP volumes
3. Collecting configurations from 7-Mode volumes, and applying the configurations to the clustered Data ONTAP volumes and SVM
4. Configuring 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. Manual: Verifying the configurations and data access in clustered Data ONTAP
6. Manual: Finish testing

The clustered Data ONTAP volumes are resynchronized.

Subproject with all volumes having LUNs (SAN only)

In precutover read-only mode, only an incremental update is performed from 7-Mode volumes to clustered Data ONTAP volumes. SAN configurations are not applied to the SVM.

Precutover read/write mode is not supported if any volume in the subproject has LUNs.

Subproject with some volumes having LUNs (SAN and NAS)

If a subproject has both SAN and NAS volumes, only NAS configurations are transitioned in the **precutover read-only** mode. Precutover read/write mode is not supported if any volume in the subproject has LUNs.

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

1. Performing an incremental update from 7-Mode volumes to clustered Data ONTAP volumes
2. The following additional steps are performed only for the NAS volumes (volumes that do not have any LUNs) in the subproject:
 - a. Breaking the SnapMirror relationship between 7-Mode volumes and clustered Data ONTAP volumes
 - b. Collecting configurations from 7-Mode volumes and applying the configurations to the clustered Data ONTAP volumes and SVM
 - c. Configuring 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.
 - d. Resynchronizing the SnapMirror relationship between 7-Mode volumes and clustered Data ONTAP volumes

Storage cutover phase

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

1. Optional: Performing an on-demand SnapMirror update to reduce the downtime after cutover
2. Manual: Disconnecting client access from the 7-Mode system
3. Performing a final SnapMirror update from 7-Mode volumes to clustered Data ONTAP volumes
4. Breaking and deleting the SnapMirror relationship between the 7-Mode volumes to clustered Data ONTAP volumes, making the clustered Data ONTAP volumes read/write
5. Applying SAN and quota configurations, if any
6. Configuring and starting the iSCSI or FC service on the SVM, if required
7. Removing 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

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

8. Optional: Taking the 7-Mode volumes offline

Post-transition steps

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

1. Perform any 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.
7-Mode Transition Tool 2.1 SAN Host Transition and Remediation Supplemental Guide
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 clustered Data ONTAP volumes.

Related tasks

Migrating data and configuration from 7-Mode volumes on page 57

How you transition volumes in a SnapMirror relationship

When you transition 7-Mode volumes that are in a SnapMirror relationship, the secondary volumes are transitioned first. After this, a volume SnapMirror relationship is established between the 7-Mode primary volumes and the clustered Data ONTAP secondary volumes. After transitioning the primary volumes, the tool establishes a volume SnapMirror relationship between the clustered Data ONTAP primary and secondary volumes.

You can perform precheck, baseline copy, incremental transfers, and precutover on the secondary and primary subprojects simultaneously; however, the storage cutover for the secondary subproject 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. Adding 7-Mode storage system and volume information
2. Gathering 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. Running the transition precheck
4. Adding cluster, SVM, and aggregate information
5. Collecting 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

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

6. Creating the data copy schedules for baseline and incremental transfers

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 clustered Data ONTAP volumes. The 7-Mode Transition Tool performs the following tasks in this phase:

1. Creating the clustered Data ONTAP volumes with read-only access
2. Setting up a transition peer relationship between the 7-Mode system and the SVM
3. Establishing a SnapMirror relationship between the 7-Mode volumes and clustered Data ONTAP volumes
4. Completing the baseline data transfer based on schedule inputs
5. Performing scheduled SnapMirror data copy updates to the clustered Data ONTAP volumes

Important: You cannot edit a subproject after starting the data copy operations.

Precutover phase

In this phase, configurations from the 7-Mode volumes are copied to the clustered Data ONTAP volumes. There are two modes for the precutover phase: precutover read-only and precutover read/write. It is best that you create separate subprojects for NAS and SAN volumes because precutover read/write mode is disabled if you have at least one SAN volume in your subproject.

Subproject with volumes that do not have LUNs (NAS only)

If a subproject does not have any volumes with LUNs, both the precutover read-only and precutover read/write modes are supported.

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

1. Performing an incremental update from 7-Mode volumes to clustered Data ONTAP volumes
2. Breaking the SnapMirror relationship between 7-Mode volumes and clustered Data ONTAP volumes
3. Collecting configurations from 7-Mode volumes and applying the configurations to the clustered Data ONTAP volumes and SVM
4. Configuring 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. Resynchronizing the SnapMirror relationship between 7-Mode volumes and clustered Data ONTAP volumes

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

1. Performing an incremental update from 7-Mode volumes to clustered Data ONTAP volumes
2. Breaking the SnapMirror relationship between 7-Mode volumes and clustered Data ONTAP volumes
3. Collecting configurations from 7-Mode volumes and applying the configurations to the clustered Data ONTAP volumes and SVM
4. Configuring 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. Manual: Verifying the configurations and data access in clustered Data ONTAP
 6. Manual: Finish testing
The clustered Data ONTAP volumes are resynchronized.

Subproject with all volumes having LUNs (SAN only)

In the precutover read-only mode, only an incremental update is performed from the 7-Mode volumes to clustered Data ONTAP volumes. SAN configurations are not applied to the SVM

Precutover read/write mode is not supported if any volume in the subproject has LUNs.

Subproject with some volumes having LUNs (SAN and NAS)

If a subproject has both SAN and NAS volumes, only NAS configurations are transitioned in the **precutover read-only** mode. Precutover read/write mode is not supported if any volume in the subproject has LUNs.

The tool performs the following steps:

1. Performing an incremental update from 7-Mode volumes to clustered Data ONTAP volumes
2. The following steps are performed only for the NAS volumes (volumes that do not have any LUNs) in the subproject:
 - a. Breaking the SnapMirror relationship between 7-Mode volumes and clustered Data ONTAP volumes
 - b. Collecting configurations from 7-Mode volumes, and applying the configurations to the clustered Data ONTAP volumes and SVM
 - c. Configuring 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.
 - d. Resynchronizing the SnapMirror relationship between 7-Mode volumes and clustered Data ONTAP volumes

Storage cutover (secondary volumes) phase

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


1. Optional: Performing an on-demand SnapMirror update on the clustered Data ONTAP secondary volumes
2. Manual: Disconnecting client access, if required
3. Performing a final SnapMirror update from the 7-Mode secondary volume to the clustered Data ONTAP secondary volume
4. Breaking and deleting the SnapMirror relationship between the 7-Mode secondary volume and the clustered Data ONTAP secondary volume, and making the destination volumes read/write
5. Applying SAN and quota configurations, if any

6. Establishing a SnapMirror relationship between the volumes on the 7-Mode primary system and the clustered Data 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 clustered Data ONTAP secondary volumes.
7. Configuring and starting the iSCSI or FC service on the SVM, if required
8. Removing 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
Note: SAN LIFs are not transitioned by the 7-Mode Transition Tool.

9. Optional: Taking the 7-Mode volumes offline

Note: The igroup to LUN mapping configuration on the 7-Mode secondary is not applied to the clustered Data ONTAP secondary volume. You must map the secondary LUNs manually after completing the storage cutover of the primary volumes.

The following illustration depicts the transition of a secondary volume:



Storage cutover (primary volumes) phase

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


1. Optional: Performing an on-demand SnapMirror update on the clustered Data ONTAP secondary volumes
2. Manual: Disconnecting client access from the 7-Mode system
3. Performing a final incremental update from the 7-Mode primary volume and the clustered Data ONTAP primary volume
4. Breaking and deleting the SnapMirror relationship between the 7-Mode primary volume and the clustered Data ONTAP primary volume, and making the destination volumes read/write
5. Applying SAN and quota configurations, if any

6. Breaking and deleting the SnapMirror relationship between the 7-Mode primary volume and the clustered Data ONTAP secondary volume
7. Setting up cluster peer and SVM peer relationships between the primary and secondary clusters
8. Setting up a SnapMirror relationship between the primary and secondary clustered Data ONTAP volumes
9. Resynchronizing the SnapMirror relationship between the clustered Data ONTAP volumes
10. Configuring and starting the iSCSI or FC service on the SVM, if required
11. Removing 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

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

12. Optional: Taking the 7-Mode volumes offline

The following illustration depicts the transition of a primary volume:



Post-transition steps

After the cutover phase finishes successfully and 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. Map the secondary LUNs to their igroups manually.
3. For SAN transition, manually reconfigure the hosts.
7-Mode Transition Tool 2.1 SAN Host Transition and Remediation Supplemental Guide
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 clustered Data ONTAP volumes.

Related tasks

[Migrating data and configuration from 7-Mode volumes](#) on page 57

NFS migration workflow

The NFS migration workflow involves adding the NFS license, creating and running a subproject by using the 7-Mode Transition Tool, and performing some manual tasks in the precutover and storage cutover phases.




Related tasks

[Migrating data and configuration from 7-Mode volumes](#) on page 57

CIFS migration workflow

The CIFS migration workflow involves creating and running a project by using the 7-Mode Transition Tool and performing some manual tasks, such as licensing CIFS, configuring a CIFS server, and verifying the configurations in the SVM before storage cutover.


**Related tasks**

[Migrating data and configuration from 7-Mode volumes](#) on page 57

FC migration workflow

The FC migration workflow involves adding the FC license, creating an FC LIF, and creating and running a subproject by using the 7-Mode Transition Tool. SAN configurations are applied only after

storage cutover. You must then perform some post-transition tasks for FC zoning and reconfigure the hosts before enabling access to the clustered Data ONTAP volumes.




Related information

[7-Mode Transition Tool 2.1 SAN Host Transition and Remediation Supplemental Guide](#)

iSCSI migration workflow

The iSCSI migration workflow involves adding the iSCSI license, creating an iSCSI LIF, and creating and running a subproject by using the 7-Mode Transition Tool. SAN configurations are

applied only after storage cutover. You must perform some post-transition tasks on the hosts before enabling access to the clustered Data ONTAP volumes.



Related information

[7-Mode Transition Tool 2.1 SAN Host Transition and Remediation Supplemental Guide](#)

Preparing for migration

Before initiating a data copy operation from 7-Mode to clustered Data 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 you migrate the 7-Mode volumes and configurations:

- The 7-Mode and clustered Data 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 references

[Transition preparation checklist](#) on page 58

Data ONTAP version requirements for migration

You should be aware of the versions of Data ONTAP operating in 7-Mode that are supported for migrating 7-Mode volumes to clustered Data ONTAP 8.2.x and 8.3.x. NAS transition to clustered Data ONTAP 8.2.x or 8.3.x is supported. However, SAN transition is supported only to clustered Data ONTAP 8.3.x.

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

License requirements for transition

Before you transition a volume from 7-Mode to clustered Data ONTAP, you must ensure that SnapMirror is licensed on the 7-Mode storage system. Additionally, SnapMirror licenses are required on the primary and secondary clusters to transition a 7-Mode volume SnapMirror relationship. The 7-Mode Transition Tool does not require any license.

If SnapMirror is already licensed on your 7-Mode system, you can use the same license for transition. If you do not have the 7-Mode SnapMirror license, you can obtain a temporary SnapMirror license for transition from your sales representative.

Feature licenses that are enabled on the 7-Mode system must be added to the cluster. For information about obtaining feature licenses on the cluster, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Related information

[Documentation on the NetApp Support Site: mysupport.netapp.com](https://mysupport.netapp.com)

7-Mode Transition Tool service must be running

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 where the tool is installed. If you access the web interface of the tool from a browser on a different system, you can shut down that system.

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 ensure 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 clustered Data ONTAP secondary volume.

- You must not initiate transition while the aggregates on either the 7-Mode system or cluster are upgrading from 32-bit to 64-bit format; otherwise transition fails.
- The 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 sessions for each secondary volume, complete the transition of the secondary sessions, and then create and complete the transition of the primary volume.

- Transition operations such as precheck and data copy cannot be run for two subprojects simultaneously if the subprojects share the same 7-Mode or the clustered Data ONTAP systems.
- 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 tool does not transition a vFiler unit as a single entity.
However, you can transition all the volumes in a vFiler unit by selecting them as a part of one or more subprojects.
- The tool does not transition root volume of a vFiler unit if the root volume is based on a qtree that belongs to the default vFiler unit.
- The 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 cluster for transition

You must have set up the cluster before transitioning a 7-Mode system and you must ensure that the cluster meets requirements such as allowing HTTPS, setting up intercluster LIFs, and verifying network connectivity for transition.

Before you begin

- The cluster and the SVM must already be set up.
For information about setting up the cluster and the SVM, see the *Clustered Data ONTAP Software Setup Guide*.
- The cluster must be healthy and none of the nodes must be in takeover mode.
- The aggregates that are used to place the transitioned volumes in the SVM 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 cluster running Data ONTAP 8.2.x, 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. For more information about aggregates, see the *Clustered Data ONTAP Physical Storage Management Guide*.
- 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 already peered SVMs.

Steps

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

```
ssh username@cluster_mgmt_IP
```

2. Enable SSLv3 or FIPS on the cluster:

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

When FIPS 140-2 compliance is enabled, SSLv3 is disabled. Data 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.

Important: 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. Use the `network interface show` command to view the firewall policy for the cluster-management LIF.

Example

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

- b. Use the `system services firewall policy show` command to verify that the firewall policy associated with the cluster-management LIF allows HTTPS access.

Example

```
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.
```

For more information about firewall policies, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

4. Create the intercluster LIF on each node of the cluster for communication between the cluster and 7-Mode system:
 - a. Use the `network interface create` command to create an intercluster LIF.

Example

```
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. If you are transitioning to clustered Data ONTAP 8.2.x, use the `network routing-groups route create` command to create a static route for the intercluster LIF.

Example

```
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
```

- c. If you are transitioning to clustered Data ONTAP 8.3.x, use the `network route create` command to create a static route for the intercluster LIF.

Example

```
cluster1::> network route create -vserver vs0 -destination
0.0.0.0/0 -gateway 10.61.208.1
```

- d. Use the `network ping` command to verify that you can use the intercluster LIF to ping the 7-Mode system.

Example

```
cluster1::> network ping -lif intercluster_lif -lif-owner
cluster1-01 -destination system7mode
system7mode is alive
```

For multipathing, you must have two intercluster LIFs on each node. For more information about network interfaces, see the *Clustered Data ONTAP Network Management Guide*.

Related information

[NetApp Documentation: Product Library A-Z](#)

Preparing the 7-Mode system for transition

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

Before you begin

All the 7-Mode volumes to be transitioned must be online.

About this task

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

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 xxxxxxxx
```

xxxxxxx is the license code you purchased.
 - b. Enable SnapMirror:


```
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` is **off**, add the SVM name to the `/etc/snapmirror.allow` file.

- If the `snapmirror.access` option is **legacy** and the `snapmirror.checkip.enable` is **on**, add the IP addresses of the intercluster LIFs to the `/etc/snapmirror.allow` file.
3. Enable HTTPS on the 7-Mode system:
options httpd.admin.ssl.enable on
 4. Enable TLS on the 7-Mode system:
options tls.enable on
Important: The best practice is to enable TLS because of the security vulnerabilities in SSLv3.
 5. 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
Important: The best practice is to disable SSLv2 and SSLv3 to avoid security vulnerabilities.
 6. 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.snapmirror.blocked ""
 - b. If you are running Data ONTAP versions 7.3.7, 8.0.3, or 8.1 and 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
 7. 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

Related information

[Documentation on the NetApp Support Site: mysupport.netapp.com](https://mysupport.netapp.com)

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 for allowing communication with the 7-Mode Transition Tool.

The following ports must open on the 7-Mode system for enabling communication with the cluster and 7-Mode Transition Tool:

Ports in the 7-Mode system	To communicate with the cluster (for SnapMirror)	To communicate with 7-Mode Transition Tool (for HTTP/HTTPS)
	• 10565/TCP • 10566/TCP • 10567/TCP • 10568/TCP • 10569/TCP • 10670/TCP	• 80/TCP • 443/TCP

The following ports must be open on the cluster for enabling communication with the 7-Mode system and 7-Mode Transition Tool:

Ports in the cluster	To communicate with the 7-Mode system (for SnapMirror)	To communicate with 7-Mode Transition Tool (for HTTP/HTTPS)
	• 10565/TCP • 10566/TCP • 10567/TCP • 10568/TCP • 10569/TCP • 10670/TCP • 11105/TCP	• 80/TCP • 443/TCP

Prerequisites for transitioning NFS configurations

You should meet certain prerequisites on the 7-Mode system and the cluster for NFS configurations to be transitioned by the 7-Mode Transition Tool. If any of the conditions is not met, the tool does not transition the configuration.

- On the 7-Mode system, the following conditions must be met:
 - NFS must be licensed.
 - If the MultiStore license is enabled, NFS must be added to the list of allowed protocols for the vFiler unit that owns the transitioning volumes.

- For transitioning a MS AD based Kerberos server to a new SVM, a DNS entry must exist for the AD domain.

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

- On the cluster, the following conditions must be met:
 - NFS must be licensed.
 - NFS must be added to the list of allowed protocols for the SVM.
 - For transitioning an MS 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 clustered Data ONTAP system must be less than or equal to five minutes.

Related concepts

[NFS migration workflow](#) on page 39

[How NFS exports are transitioned](#) on page 100

[NFS only configuration transition](#) on page 99

Prerequisites for transitioning CIFS configurations

You should meet certain prerequisites on the 7-Mode system and on the cluster for CIFS configurations to be transitioned by the 7-Mode Transition Tool. If any of the condition is not met, the tool does not transition the configuration.

On the 7-Mode system, the following conditions must be met:

- CIFS license must be added.
- 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 must be set up and running.
- The authentication type for CIFS must be Active Directory (AD).

On the cluster, the following conditions must be met:

- 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.
 - CIFS server must be manually configured before precutover.
- You can create the CIFS server on the SVM in the following two ways:

If you want to..	Conditions
Transfer or preserve CIFS server identity to the target SVM	◦ You must plan to transition all the volumes in source 7-Mode system or vFiler unit in a single subproject. This 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 a given subproject 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 subproject. ◦ Before precutover, reconfigure the CIFS server on the 7-Mode by using a temporary CIFS identity. This allows the original CIFS server identity to be configured on the SVM. Ensure that CIFS server is running on the 7-Mode during the precutover operation with the new temporary identity. This is required to read CIFS configurations from 7-Mode during precutover. ◦ Configure the CIFS server on the target SVM with the original 7-Mode CIFS identity. After ensuring these conditions, you can perform the precutover operation. You must then plan to perform the storage cutover immediately after precutover for enabling client access to the clustered Data ONTAP volumes.
Use a new identity	◦ Before precutover, configure CIFS server on the target SVM with a new CIFS identity. ◦ Ensure that CIFS server is up and running on the 7-Mode system during the precutover operation. This is required to read CIFS configurations from 7-Mode during the precutover operation. After ensuring these conditions, you can perform the precutover operation. You can then test the SVM configurations and then plan to perform the storage cutover.

Related concepts

[CIFS migration workflow](#) on page 40

[Considerations for transitioning CIFS local users and groups](#) on page 103

Prerequisites and considerations for transitioning 7-Mode volumes to a MetroCluster configuration in clustered Data ONTAP

You must be aware of the requirements and considerations for transitioning 7-Mode volumes to a MetroCluster configuration in clustered Data ONTAP.

Prerequisites

- The MetroCluster configuration in clustered Data 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 giveback.
- The nodes in the MetroCluster configuration in clustered Data ONTAP must not be switched over or waiting for switchback.

Considerations

- You can transition volumes from a 7-Mode controller, HA configuration, or a MetroCluster configuration to a MetroCluster configuration in clustered Data 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.
- 7-Mode volumes on unmirrored aggregates can be transitioned only to mirrored aggregates.
- Different subnets configured for a 7-Mode fabric MetroCluster configuration cannot be configured on the MetroCluster configuration in clustered Data ONTAP.
- The preferred port configured in a 7-Mode fabric MetroCluster configuration cannot be configured for the MetroCluster configurations in clustered Data 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 clustered Data ONTAP. It is best to share the switch fabrics only for the duration of the transition.
[Clustered Data ONTAP 8.3 MetroCluster Installation and Configuration Guide](#)
- 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 post-transition.

Related concepts

[Impact of takeover and giveback on transition](#) on page 106

Related tasks

[Configuring cron job schedules on the remote site after transitioning a MetroCluster configuration](#) on page 111

[Transitioning a MetroCluster configuration that failed due to switchover or switchback](#) on page 116

Preparation for SAN transition

You should use OnCommand System Manager for provisioning SVM and configuring the data LIFs on the SVM.


Preparing for transition to clustered Data ONTAP 8.3 and later

32-bit data is not supported in clustered Data ONTAP 8.3 and later. If your 7-Mode system has 32-bit data, you must expand the 32-bit aggregates to 64-bit, and then find and remove any 32-bit data.

Because all 7-Mode versions do not support the capability of expanding 32-bit aggregates and removing 32-bit data, you might have to upgrade your 7-Mode system before transition.

About this task

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



Steps

1. [Expanding an aggregate to the 64-bit format](#) on page 54
2. [Finding and removing 32-bit data from source volumes and Snapshot copies](#) on page 55

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.

Before you begin

- 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.

About this task

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. Take 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	Take 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. Optional: 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
```

Result

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

Finding and removing 32-bit data from source 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.

Before you begin

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

About this task

You must repeat the steps in this task for each aggregate that contains 32-bit data.

Steps

- Enter advanced mode:

```
priv set advanced
```

- 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.

- 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.

- Display the format of all Snapshot copies on the system:

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

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

Attention: 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
```

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 create a project and configure its subprojects, and then start and complete all the subprojects created in that project.


Before you begin

- You must ensure that the 7-Mode controllers and clusters that you want to include in the transition are 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 was installed. The service is set to automatic by default, and should start when you restart the machine.

About this task

- You should not perform both 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, as it might cause undesired results.

The following image illustrates the migration process:



Steps

1. [Transition preparation checklist](#) on page 58
2. [Adding controllers and clusters](#) on page 62
3. [Creating a transition project](#) on page 63
4. [Running prechecks](#) on page 66
5. [Starting baseline data copy](#) on page 68
6. [Applying 7-Mode configurations](#) on page 69
7. [Completing a transition project](#) on page 72

Related concepts

[NFS migration workflow](#) on page 39

[CIFS migration workflow](#) on page 40

[How you transition a stand-alone volume](#) on page 30

[How you transition volumes in a SnapMirror relationship](#) on page 34

Transition preparation checklist

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

- [Data ONTAP version requirements](#) on page 58
- [Licensing requirements](#) on page 58
- [SnapMirror requirements on the 7-Mode system](#) on page 59
- [Volume options on the 7-Mode system](#) on page 59
- [Managing access to the cluster](#) on page 59
- [Managing access to the 7-Mode system](#) on page 59
- [Networking requirements](#) on page 60
- [Port requirements](#) on page 60
- [NFS requirements](#) on page 61
- [CIFS requirements](#) on page 61

Data ONTAP version requirements

Item	Yes
Supported 7-Mode version NetApp Interoperability Matrix Tool	
Your cluster must be running one of the following Data ONTAP versions: • Data ONTAP 8.2.x • Data ONTAP 8.3.x	

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	
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	
options snapmirror.enable on	
options interface.snapmirror.blocked ""	
Verify if one of the following is true: - The snapmirror.access option is set to all - The snapmirror.access option is set to the IP addresses of all the intercluster LIFs on the cluster - If the snapmirror.access option is set to legacy and the snapmirror.checkip.enable option is off, the SVM name is added to the /etc/snapmirror.allow file - If the snapmirror.access option is set to legacy and the snapmirror.checkip.enable option is on, the IP addresses of the intercluster LIFs are added to the /etc/snapmirror.allow file	

Volume settings on the 7-Mode system

Item	Yes
Volume is online	
Volume is not restricted	
The following volume options are disabled: - no_i2p - read_realloc - nvfail	

Managing access to the cluster

Item	Yes
SSL is enabled system services web show	
HTTPS is allowed on the cluster-management LIF system services firewall policy show	

Managing access to the 7-Mode system

Item	Yes
HTTPS is enabled options httpd.admin.ssl.enable on	

Item	Yes
SSL is enabled secureadmin setup ssl options ssl.enable on	
SSLv2 and SSLv3 are disabled options ssl.v2.enable off options ssl.v3.enable off	

Networking requirements

Item	Yes
Cluster is reachable using the cluster-management LIF	
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 the 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 • 10565/TCP • 10566/TCP • 10567/TCP • 10568/TCP • 10569/TCP • 10670/TCP • 80/TCP • 443/TCP	

Item	Yes
Cluster	
• 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)	
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 concepts

[Preparing for migration](#) on page 43

Adding controllers and clusters

You can add or remove the 7-Mode controllers and clusters that are required for the transition.

About this task

- The 7-Mode controllers that are included for assessment are automatically added for transition.
- You cannot remove the controllers and clusters that are part of any active subproject.
- The 7-Mode controller and cluster information that you provide is not persistent.
If the Transition Tool service is restarted for any reason, 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 subproject.
- If a 7-Mode controller is part of a high availability (HA) pair, the tool does not ask for credentials of the HA partner of the 7-Mode controller (unless the HA partner is part of another active subproject.)

Steps

1. Log in to the application and click **Get Started** in the **Migrate** section.
2. From the left pane, click **Storage Systems**.
3. In the **Hostname** field, enter the fully qualified domain name (FQDN) or the IP address of the 7-Mode controller or the clustered Data ONTAP system.

For a cluster, you can specify the IP address (or FQDN) of the management interface, or the IP address (or FQDN) of any node within the cluster. 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.
4. In **Username** and **Password**, enter the administrator credentials for the specified host, and click **Add**.

The tool detects whether the added system is a 7-Mode controller or a cluster, and includes it in the appropriate table.
5. Repeat Steps 2 and 3 to add all the controllers and clusters that you require for the transition.

Result

The 7-Mode controllers and clusters are added to the 7-Mode Controllers and Clustered Data ONTAP Systems tables, respectively.

Note: The Status column indicates if the credentials of the system are missing or the credentials have changed from what was initially entered in the tool. You can click the  icon and 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, running prechecks, and creating data copy schedules for volume SnapMirror transfers.

Before you begin

You must have created the required SVM on the cluster.

About this task

All volumes within a project are migrated to the same Storage Virtual Machine (SVM). If you want to migrate the volumes to different Storage Virtual Machines (SVMs), you must create multiple projects. Such active projects cannot write to the same SVM.

You should create separate subprojects for NAS and SAN volumes because precutover read/write mode is disabled if you have at least one SAN volume in your subproject.

Steps

1. Log in to the tool and click **Get Started** in the **Migrate** section.
2. From the left pane, click **New Transition Project** to create a new transition project.
If the controller and cluster required by the new project are not added, you can enter the details in the Enter Device Credentials pane.
3. In the **New Transition Project Name** dialog box, enter a name for the project and click **Continue**.
A project name can contain up to 25 alphanumeric characters. Special characters and spaces are not allowed.
4. Select the 7-Mode volumes that you want to transition.
 - a. From the **Choose Origin Volumes to Transition** section, select the 7-Mode controller or the vFiler unit from which you want to add volumes.
 - b. Add the volumes you want to include in the project by selecting one of the following transition options:

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

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

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

Tip: It is a best practice to have all 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 volumes you want to include in the project, click **Create Project and Continue**.
5. Select the 7-Mode IP address and multipath IP address to be used for data copy.
 - a. Enter the 7-Mode data copy IP address.
By default, this is prepopulated with the management IP address of the 7-Mode system. If required, this IP address can be changed to any valid IPv4 address with data copy permission.
 - b. Optional: If you want to use multiple paths for load balancing the data transfers, enter an IP address in the **IP Configuration** pane and click **Next**.
6. Optional: Verify if the 7-Mode volumes can be transitioned by running prechecks.
 - a. From the **7-Mode Precheck** pane, click **Run Prechecks**.
The prechecks might take several minutes to complete. After the prechecks are complete, the Precheck page is displayed with the results.
 - b. To save the Precheck results, click **Save as CSV**.
 - c. Click **Next**.
7. From the **SVM Mapping** pane, select the SVM to which the 7-Mode volumes must be transitioned, and click **Next**.
For transitioning 7-Mode volumes to a MetroCluster configuration in clustered Data ONTAP, the SVM subtype must be **sync-source**.
Note: If you select an SVM that belongs to clustered Data ONTAP 8.2, a dialog box is displayed for confirming if there are local users and groups, CIFS shares or files 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 belonging to clustered Data ONTAP 8.2.1 or later.
8. Optional: 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 clustered Data ONTAP system.
9. Select the target volumes for transition to map each of the source volume to the desired aggregate.
 - a. From the **Volume Mapping** pane, select the aggregates to which the 7-Mode volumes must be copied.
 - b. Optional: To change the name of target volume on the cluster, enter a different name in **Target Volume**.
 - c. Click **Next**.

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

10. Optional: From the **Interface Mapping** pane, provide information about the LIFs that must be created on the SVM.

Note: 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	- Select the 7-Mode IP address. - Click Choose Node and Port. - From the Choose Node and Port dialog box, select the target node and port to which the LIF must be migrated, and click Save. - Click Next.
Create a new LIF	- Click New LIF. - In the Provide details for additional LIF in the destination cluster dialog box, enter the details for the new LIF, and click Add. - Click Next.

To ensure network connectivity after a successful transition, you must transition the 7-Mode IP addresses to a similar network topology in clustered Data 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 clustered Data ONTAP. Similarly, IP addresses configured on VLAN ports or interface groups should be transitioned to appropriate VLAN ports or interface groups in clustered Data ONTAP.

11. Configure the data copy schedules for baseline and incremental transfers, number of concurrent volume SnapMirror transfers, and throttle limit for the SnapMirror transfers for transition.

You can provide data copy schedules and 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 subproject. For example, you can create customized schedules for weekdays and weekends.

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

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

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

- 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 the volumes in the subproject).

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

- g. Click **Create**.

The new schedule is added to the Transition Schedule pane.

12. Verify the 7-Mode system and cluster for transition by running a final precheck.
 - a. From the **Final Precheck** pane, click **Run Prechecks**.
 - b. Review the precheck results, fix any errors manually, and then click **Next**.
13. From the **Subproject Summary** pane, save the project.

If you want to transition a subproject with...	Then...
Stand-alone volumes	Click Save and go to Dashboard .
Volumes in a SnapMirror relationship	Click Save and Configure Next Subproject to save the current subproject and configure the next subproject.

Related concepts

[Considerations for creating a data copy schedule](#) on page 79

[Managing SnapMirror transfers and schedule](#) on page 74

Related tasks

[Creating a data copy schedule for SnapMirror transfers](#) on page 74

[Transitioning volumes by excluding local user and groups or netgroup configurations](#) on page 82

[Managing logical interfaces](#) on page 83

[Removing volumes from a subproject](#) on page 83

Running prechecks

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

Before you begin

You must have created at least one data copy schedule.

About this task

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 subproject for which you want to run the prechecks.
2. Click **Run Prechecks**.

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

Note: 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. Click **View Detailed Results**.
4. 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**.
5. 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.

After you finish

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

Prechecks can be run at two different stages while creating and configuring a transition project:

- Initial prechecks
You can run the initial prechecks after selecting the 7-Mode volumes for transition.
- Final prechecks
You can run the final prechecks on the 7-Mode volumes selected for transition and the SVMs and aggregates to which the selected volumes are transitioned.

The problems 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 some 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 clustered Data 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 need to complete some manual tasks for enabling these configurations in clustered Data ONTAP. The following are some 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 clustered Data ONTAP. You can start baseline data copy operation for individual subprojects. You should stop unnecessary system processes and network activity during the data copy.

Before you begin

You must have created at least one data copy schedule.

About this task

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 clustered Data ONTAP destination controllers
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.
- 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 subproject 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 subproject.

Related tasks

[Creating a data copy schedule for SnapMirror transfers](#) on page 74

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 clustered Data ONTAP volumes. SAN configuration is not transitioned in this phase.

About this task

The configurations are applied in the precutover phase, which has two modes:

- Precutover read-only mode
- Precutover read/write mode

The precutover read/write mode is disabled if there is at least one SAN volume in the subproject.

If a subproject has 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.

Choices

- [Applying 7-Mode configurations for volumes without any LUNs \(NAS only\)](#) on page 69
- [Applying configurations on volumes that have at least one LUN](#) on page 71

Related concepts

[Considerations for quotas](#) on page 105

Related tasks

[Managing logical interfaces](#) on page 83

Applying 7-Mode configurations for volumes without any LUNs (NAS only)

If you create subprojects with only NAS volumes, you can run precutover in two modes: precutover read-only mode and precutover read/write mode. In the precutover read/write mode, you can test the data and configurations in clustered Data ONTAP.

Steps

1. On the **Dashboard**, select the subproject.
2. You can apply the configurations in either the read-only mode or read-write mode.

If you want to apply all configurations from the 7-Mode system to the clustered Data ONTAP volumes in the...	Then...
Read-only mode	a. Click Apply Configuration .
Read/write mode	a. Select the Test Mode check box. b. Click Apply Configuration . The clustered Data ONTAP volumes are made read-write and you can test the configurations and data-access operations.

3. Optional: In the **Apply Configuration (Precutover)** dialog box, select **Apply configuration in test mode**, if you want to apply configurations in the read/write mode.
4. Optional: In the **Apply Configuration (Precutover)** dialog box, 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. Optional: Enter the maximum number of concurrent SnapMirror transfers to run during transition.
 - b. Optional: Enter the percentage of the available streams that can be used for SnapMirror transfers.

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

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

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

 - a. Enter the Kerberos server details such as hostname, IP address, username, and password.

Note: To transition the Kerberos configuration, at least one LIF needs to be transitioned as part of the subproject and the LIF must be resolvable to a host name.
6. Click **Continue**.

The Operation Progress dialog box is displayed, and the copy configuration operation is started.
7. Optional: If the configuration transition is performed in the read/write mode, click **Finish Testing** after the testing and verification of the configurations is completed.

This mode should be used only for testing purposes; it is not a method to allow both modes to remain operational. 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 present for that project at that time) the clustered Data ONTAP volumes. Any data written to the 7-Mode is resynchronized with the clustered Data ONTAP volumes.

Note: For 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.

Result

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.

Applying configurations on volumes that have at least one LUN

If you have at least one LUN in the transitioning volumes, precutover read/write mode is disabled. Also, SAN configurations are not transitioned in the precutover mode, but are transitioned during storage cutover.

Steps

1. On the **Dashboard**, select the subproject.
2. Click **Apply Configuration**.
3. Optional: In the **Apply Configuration (Precutover)** dialog box, 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. Optional: Enter the maximum number of concurrent SnapMirror transfers to run during transition.
 - b. Optional: Enter the percentage of the available streams that can be used for SnapMirror transfers.
By default, the tool uses 50% of the available volume SnapMirror transfers.
 - c. Enter the throttle limit or select **Maximum** to use the maximum bandwidth.
By default, the tool uses maximum throttle for configuration transition.
4. Optional: In the **Apply Configuration (Precutover)** dialog box, select the **Transition Kerberos Configuration** check box to provide the UNIX based or Microsoft AD based Kerberos server configuration details for transition.

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

 - a. Enter the Kerberos server details such as hostname, IP address, username, and password.
Note: To transition the Kerberos configuration, at least one LIF needs to be transitioned as part of the subproject 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.

Result

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.

Note: SAN LIFs are not configured.

Completing a transition project

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

Before you begin

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.

About this task

The storage cutover finishes in only 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. It is not necessary 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.

Steps

- From **Dashboard**, select the subproject that you want to complete.
- Click **Update Now** to perform a final SnapMirror update and wait for it to complete.
It is a best practice to perform an incremental transfer as close to the storage cutover as possible. This process minimizes the time that clients have to be disconnected from the volumes.
- Manual: Disconnect client access.
- Click **Complete Transition**.
 - Optional: If you want to keep the 7-Mode source volumes online after the transition, clear the **Take source volumes offline after transition** check box.
By default, this option is selected and the source volumes are taken offline.

- b. Optional: 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 message is displayed, warning you that the cutover operation is disruptive.

- c. Click **Yes** to continue.

After the cutover operation finishes successfully, a confirmation message is displayed.

- d. Click **OK**.

Result

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.

After you finish

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

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

Related concepts

[*Considerations for transition*](#) on page 98

Managing a transition project

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

Editing a subproject

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

Before you begin

You must have created a subproject, but not started it.

About this task

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

Steps

1. From **Dashboard**, select the subproject that you want to edit.
2. Click **Edit Configuration**.
The Volume selection pane of the New Transition Project wizard is displayed.
3. Complete the wizard to edit your subproject.

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 tasks

[Creating a transition project](#) on page 63

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 subproject. 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.

Before you begin

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

About this task

- 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 subproject.
- You can create a maximum of 7 schedules per subproject; however, the schedules cannot overlap within a subproject.
For example, you can create customized schedules for business hours and non-business hours, DR hours and non-DR hours, weekdays and weekends.

Note: It is recommended that the data copy schedules do not overlap across different subprojects, if the subprojects use the same 7-Mode controller or cluster.

- 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.

Note: 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. You can create schedules from either the **Data Copy Schedule** page or the **Dashboard**.

If you want to...	Then...
Create a schedule from the Data Copy Schedule page	Click Create Schedule in the Data Copy Schedule pane.
Create a schedule from the Dashboard	a. Select the subproject. b. Click the Data Copy Schedule tab. c. Click Create Schedule.

The Create Data Copy Schedule dialog box is displayed.

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...	Then...
Run the data copy operations daily	Select Daily . This is the default value.
Run the data copy operations only on certain days	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.
Note: The duration of a schedule should not exceed a week (167 hours and 30 minutes).

Example

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 minute
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.
Note: 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.

Example

Consider your platform supports a maximum of 100 concurrent volume SnapMirror transfers and only 60 are currently available. And, you have specified the following values:

- Percentage of the available volume SnapMirror transfers option as 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 as 25

The tool sets the maximum number of concurrent volume SnapMirror transfers to 25, which is the lowest of the two values.

- b. 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 subproject.

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

Example

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.

Result

The schedules become effective only when the subproject is in data copy or 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 am to 5 pm	50% of available transfers
Monday to Friday	11.30 pm to 2.30 am	75 transfers used for DR
Monday to Friday	2.30 am to 9 am and 5 pm to 11.30 pm	25% of available transfers
Saturday to Monday	2.30 am (Saturday) to 9 am (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	Mon-Fri
	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	Mon-Fri
	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
non_peak_non_dr1	Days Range	Mon-Fri
	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	Mon-Fri
	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	Sat
	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.

Before you begin

The subproject must be in the preparation, data copy, or precutover phase.

About this task

- 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 subproject.

Steps

1. From the **Dashboard**, select a subproject, and click the **Data Copy Schedule** tab.

All the existing schedules for the subproject are displayed.

You can also edit or delete schedules from the Data Copy Schedule page.

2. Depending on whether you want to edit or delete a schedule, perform the following steps:

If you want to...	Then...
Edit an existing schedule	a. Click  . The Modify Schedule dialog box is displayed. Note: You can edit the schedule and SnapMirror parameters for the data copy operations. b. Make the required changes, and click Save.
Delete a schedule	a. Click  . The schedule is deleted from the pane. Note: At least one schedule is required for data transfers, and therefore you should not delete all the schedules.

Example

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

Schedule	Number of volumes and data copy state	Maximum number of concurrent SnapMirror transfers	Throttle limit	Throttle used by each transfer
Before modifying	5 volumes waiting for baseline	5	500 MBps	100 MBps
After modifying	- For 2 volumes, baseline is in progress using 100 MBps throttle - For 3 volumes, baseline is completed and are waiting for updates	5	250 MBps	2 volume which are in baseline continue to use the 100 MBps throttle 3 volumes for which the baseline is completed use the modified 50 MBps throttle for the updates

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

Performing on-demand SnapMirror updates

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

About this task

- 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.

Steps

- Optional: 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.
- Optional: 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.
 - Enter the maximum number of concurrent SnapMirror transfers to run during transition.
 - 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 tasks

[Starting baseline data copy](#) on page 68

[Creating a data copy schedule for SnapMirror transfers](#) on page 74

Transitioning volumes by excluding local user and groups or netgroup configurations

By default, all 7-Mode configurations are transitioned to clustered Data ONTAP. You can choose to exclude CIFS local users and groups, UNIX user and groups, or netgroup configurations from transition by using the 7-Mode Transition Tool CLI.

Before you begin

You must have created a project by using the 7-Mode Transition Tool web interface.

About this task

With the exception of this situation where you might want to exclude certain 7-Mode configurations during transition, you must use only one interface for an end-to-end transition workflow to avoid any undesirable configurations.

Steps

1. Depending on the configurations that you want to exclude, do one or more of the following:

If you want to exclude transitioning...	Enter the following command...
CIFS local users and groups	transition property-set -s <i>sub-project-name</i> -p ignore-local-users-groups-transition -v true
UNIX users and groups	transition property-set -s <i>sub-project-name</i> -p ignore-unix-users-groups-transition -v true
Netgroups	transition property-set -s <i>sub-project-name</i> -p ignore-netgroups-transition -v true

2. Verify the value of the property set for excluding the configuration by using the following command:

transition property-get -s *sub-project-name* -p *property_name*

property_name can be ignore-local-users-groups-transition, ignore-unix-users-groups-transition, or ignore-netgroups-transition depending on the configuration that you chose to exclude from transition.

After you finish

Perform data copy operations and complete the transition of the subproject from the 7-Mode Transition Tool web interface.

Managing logical interfaces

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

About this task

The LIFs are configured on the Storage Virtual Machine (SVM) in the precutover phase:

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

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

Steps

1. Select a subproject from **Dashboard**.
2. To edit an existing logical interface, perform the following steps:

If you want to...	Then...
Edit an existing logical interface	a. Click the Edit icon in the table. The Modify LIF dialog box is displayed. b. Make the required changes and click Save.
Remove an existing logical interface	a. Click  in the table. Note: You can remove a logical interface only when the subproject is in preparation, baseline data copy, pause, precutover, or update state. However, you cannot remove a logical interface if the cutover operation has failed.
Add a new logical interface	a. Click Add LIF below the table. The Add LIF dialog box is displayed. b. Select either Add New LIF (to add a new network interface) or Transition existing 7-mode IP (to use an existing network interface). c. Enter the required values. d. Click Save to add the new interface.

Removing volumes from a subproject

You can remove volumes from a subproject during the transition process; for example, if the subproject has failed due to a few volumes.

About this task

The volumes can be in any of the following state: offline or restricted.

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

Steps

1. From **Dashboard** select a subproject, and click the **Volumes** tab.

All the existing volumes for the subproject are displayed.

2. Click .

The volume is removed from the subproject.

After you finish

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

1. From the cluster, delete the clustered Data ONTAP volumes that were created as part of transition.
[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)
2. From the 7-Mode system, release the SnapMirror relationships to the clustered Data 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 clustered Data ONTAP secondary volumes.
 - b. From the primary 7-Mode system, release the SnapMirror relationships to the clustered Data ONTAP secondary volumes.

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

Pausing and resuming a subproject

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

About this task

When you pause a primary subproject, 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 clustered Data ONTAP secondary volume continues to run according to the schedule. This ensures continued data protection for the primary volumes.

Steps

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

The data copy operation for all the volumes in the subproject 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. Resume the paused subproject by clicking **Resume**.

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

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

Related tasks

[Starting baseline data copy](#) on page 68

Aborting a subproject

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

Before you begin

The baseline copy must have been started for the subproject.

About this task

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

- If both subprojects have been started, aborting one subproject aborts the associated subproject. For example, if you abort a primary subproject, the secondary subproject is also aborted.
- If the abort operation fails on the current subproject, the associated subproject is not aborted.
- If only one active subproject is aborted, the start operation on the associated subproject fails.
- When you abort a primary subproject, the copy operation from the 7-Mode primary volume to the clustered Data ONTAP secondary volume is not aborted.
Only the copy operation from the 7-Mode primary volume to the corresponding clustered Data ONTAP primary volume is aborted.

Steps

1. From the **Dashboard**, select the subproject that you want to abort.
2. Click **Abort**.

After you finish

1. From the cluster, delete the clustered Data ONTAP volumes that were created as part of transition.
2. From the 7-Mode system, release the SnapMirror relationships to the clustered Data ONTAP volumes.
3. 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 clustered Data ONTAP secondary volumes.
 - b. From the primary 7-Mode system, release the SnapMirror relationships to the clustered Data ONTAP secondary volumes.

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

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

Deleting a project

You can delete a project from Dashboard.

Before you begin

- You must have created a project.

Steps

1. From **Dashboard**, select the project that you want to delete.
2. Click **Delete Project**.

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

You cannot delete a project if transition is in progress for any of the subproject contained in it.

Transitioning volumes using the 7-Mode Transition Tool CLI

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

Planning a session 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 sessions in your environment.

Bidirectional SnapMirror relationship

To transition a bidirectional SnapMirror relationship, you must create different transition sessions 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 session. Similarly, you cannot combine the volumes vf2:vol1 and vf2:vol2 in one transition session. You must create separate transition sessions for each volume to transition the volume SnapMirror relationship.

Volumes with CIFS configuration

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

For example, if 10 volumes in a 7-Mode system or a vFile unit have associated CIFS shares, home-directory search path, and audit configuration, these 10 volumes must be transitioned in one session.

This ensures that all the volumes and CIFS configuration are completely applied on the Vserver 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 session to group all the secondary volumes and complete the transition of all secondary volumes in that session. You can then create a primary session to transition the primary volume and complete the SnapMirror transition.

Transitioning volumes by using the 7-Mode Transition Tool command-line interface

You can create transition sessions and perform the transition in stages by using the 7-Mode Transition Tool command-line interface.

Steps

1. [Creating a transition session](#) on page 88
2. [Running prechecks](#) on page 91
3. [Starting data copy for a transition session](#) on page 92
4. [Applying configurations to the clustered Data ONTAP volumes](#) on page 93
5. [Completing the transition](#) on page 94

Creating a transition session

You can create a transition session that identifies the 7-Mode objects to be transitioned and how they map to the clustered Data ONTAP objects. While creating the session, the 7-Mode Transition Tool collects and saves the information about the transition objects. You can use this session information to perform the transition in stages.

Before you begin

You must have gathered the following information:

- Session 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) 2.1** 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 session:
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.

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

Example

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

3. Create a transition session:

```
transition create -s session_name -t session_type -n 7-Mode_system -c
data-copy-ipaddress [-f vfiler_name] [-h cluster -v vsver_name]
```

session_name is the name of the transition session.

session_type is the session 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.

Note: The host names of your 7-Mode system and cluster should be the same as those you had provided while creating the subproject.

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

vfiler_name is the name of the non-default vFiler unit.

vsver_name is the name of the SVM.

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

Example

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

The transition session moves to the preparation state.

4. Add the 7-Mode volume and clustered Data ONTAP volume to the session:

```
transition volumepair add -s session_name -v 7-Mode_volume -c
vsver_volume -g aggr_name
```

7-Mode_volume is the 7-Mode volume name.

vsver_volume is the clustered Data ONTAP volume on the SVM.

Note: You can also add a clustered Data 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 clustered Data ONTAP volume *vsver_volume* is created during the start operation.

Example

```
7-Mode Transition Tool>transition volumepair add -s sample_session -v
vol1 -c vol1 -g aggr1
```

5. Optional: Add the IP addresses to be transitioned to the session:

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

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

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

Example

```
7-Mode Transition Tool>transition lif add -s sample_session -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 clustered Data ONTAP system in the **up** state during precutover.

6. Optional: For transitioning a volume SnapMirror relationship, add details about the primary system while transitioning the secondary volumes and secondary system when transitioning the primary volumes:

If you want to add information about the...	Enter...
Primary system (when transitioning the secondary volumes)	transition add-primary-seven-mode-system -s session_name -h source-host -f management-ipaddress -d data-copy-ipaddress [-m multipathing-ipaddress]
Secondary system (when transitioning the primary volumes)	transition add-secondary-cluster-mode-system -s session_name -h c-mode-host-name -v vserver-name

session_name is the name of the transition session.

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 session have been transitioned.

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

7. Create a data copy schedule:

```
transition schedule add -s session_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 session-snapmirror-throttle
```

Example

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 -s sample_session -n dr_active -d 1-5 -b 23:30 -e 03:00 -c 25 -x 200 -u 00:30
```

- Optional: View the detailed information about the transition session created:

```
transition show -s session-name
```

Running prechecks

After creating a transition session, you can validate the session to verify if 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.

Before you begin

The transition session must be in the preparation state.

About this task

The `transition 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.

Step

- Verify the compatibility for transition by entering the following command:

```
transition precheck -s session_name
```

Example

```
7-Mode Transition Tool>transition precheck -s session_finance

[1/14 ] Session status checks
[   Errors   ]

Operation summary:
-----
The 'precheck' operation on the session 'session_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 session 'session_finance'.
```

After you finish

You must resolve all 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 session

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

Before you begin

- The session must be in the preparation state.
- You must have run a precheck on the session and resolved all errors.
- You must have configured a dedicated, high-bandwidth, low-latency network between the source and destination systems for data transfer.

Note: 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.

About this task

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 clustered Data ONTAP volumes.

Step

1. To start the data copy, enter the following command:

```
transition start -s session_name
```

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

Example

```
7-Mode Transition Tool>transition start -s session_finance

[1/17 ] Session 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 session moves to the copy-baseline state. After the baseline transfer is completed, the session moves to the copy-update state.

Applying configurations to the clustered Data 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 clustered Data ONTAP volumes. SAN configurations are not transitioned in this phase.

About this task

The configurations are applied in the precutover phase, which has two modes:

- Precutover read-only mode
- Precutover read/write mode

Step

1. Depending on the mode in which you want to run the precutover operation, do one of the following:

If you want to run the transition operation in the...	Perform these steps...
Read-only mode	transition precutover -s <i>session_name</i> -m ro_test -c max_cap_concurrent_transfers -b snapmirror_throttle
Read/write mode	- Apply the configurations to clustered Data ONTAP by using the following command: transition precutover -s <i>session_name</i> -m rw_test -c max_cap_concurrent_transfers -b snapmirror_throttle - Test and verify the configurations in clustered Data ONTAP manually. - Start the resynchronization operation by using the following command: transition resync -s <i>session_name</i>

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 session must be completed before completing the transition of the primary session.

Steps

1. Perform a data copy from the 7-Mode volumes to the clustered Data ONTAP volumes by using the following command:

```
transition update -s session_name -t available-transfers-percentage -c max-cap-concurrent-transfers -x snapmirror-throttle -r interactive
```

This helps in reducing the down time during storage cutover.

2. Manual: Disconnect client access.
3. Complete the transition process by entering the following command:

```
transition cutover -s session_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**.

Note: 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.

After you finish

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

Commands for managing transition

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

- [Commands to manage transition operations](#) on page 95
- [Commands to manage volume SnapMirror relationships](#) on page 95
- [Commands to manage credentials of Data ONTAP systems](#) on page 96
- [Commands to manage the volume pairs in a session](#) on page 96
- [Commands to manage LIF details](#) on page 96
- [Commands to manage session properties](#) on page 96
- [Commands to manage transition jobs](#) on page 97
- [Commands to manage transition schedules](#) on page 97
- [Command to collect tool logs](#) on page 97

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 menu` commands.

Commands to manage transition operations

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

If you want to...	Use this command...
Remove information about a clustered Data ONTAP secondary system from a primary session	<code>transition remove-secondary-cluster-mode-system</code>
List the 7-Mode primary systems added to a secondary session	<code>transition show-primary-seven-mode-systems</code>
List the clustered Data ONTAP secondary systems added to a primary session	<code>transition show-secondary-cluster-mode-systems</code>

Commands to manage credentials of Data ONTAP systems

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

Commands to manage the volume pairs in a session

If you want to...	Use this command...
Add the 7-Mode volume and the clustered Data ONTAP volume (volume pair) to a session	<code>transition volumepair add</code>
Modify the clustered Data ONTAP volume of a session	<code>transition volumepair modify</code>
Remove a volume pair from a session	<code>transition volumepair remove</code>
List the volume pairs in a session	<code>transition volumepair show</code>

Commands to manage LIF details

If you want to...	Use this command...
Add an SVM data LIF to a session	<code>transition lif add</code>
Modify an SVM data LIF of a session	<code>transition lif modify</code>
Remove LIFs from a session	<code>transition lif remove</code>
List all the LIFs added to a session	<code>transition lif show</code>

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

Commands to manage session properties

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

If you want to...	Use this command...
Get the value of a specific session property	<code>transition 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 session 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 schedule add</code>
Modify a SnapMirror schedule of the session	<code>transition schedule modify</code>
Remove the SnapMirror schedules from the session	<code>transition schedule remove</code>
List all the SnapMirror schedules in a session	<code>transition schedule show</code>

Command to collect tool logs

If you want to...	Use this command...
Collects the log files of the 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 tasks

[Migrating data and configuration from 7-Mode volumes](#) on page 57

Considerations for transition

You must be aware about the behavior of certain features, such as CIFS local users and groups and NFS exports during transition. You might need to perform some manual tasks before clients can be reconnected to access the clustered Data ONTAP volumes.

Related tasks

[Migrating data and configuration from 7-Mode volumes](#) on page 57

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 clustered Data ONTAP systems depends 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) are based on the schedules you configure while creating the subproject.
- If the subprojects use the same 7-Mode controller or cluster, you should ensure that the data copy schedules do not overlap across different subprojects.
- You can ensure that your existing backup and disaster recovery (DR) operations are not impacted by 7-Mode Transition Tool transition operations in the following ways:
 - You should create SnapMirror data copy schedules for a subproject 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 subproject.
- For secondary subprojects, after cutover, the incremental SnapMirror updates for the SnapMirror relationship between the 7-Mode primary volumes and the clustered Data 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 concepts

[Considerations for creating a data copy schedule](#) on page 79

Related tasks

[Creating a data copy schedule for SnapMirror transfers](#) on page 74

Related information

[Data ONTAP 8.2 Data Protection Online Backup and Recovery Guide for 7-Mode Documentation on the NetApp Support Site: \[mysupport.netapp.com\]\(http://mysupport.netapp.com\)](#)

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 clustered Data 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 tasks

[Performing on-demand SnapMirror updates](#) on page 81

NFS only configuration transition

NFS configuration is transitioned (transition of CIFS configuration is ignored) when CIFS is configured and running on the 7-Mode storage system.

If all the volumes and qtrees selected for transition are not configured for serving CIFS requests, (that is, security style of all the volumes and qtrees is UNIX-based and there are no CIFS shares or home-directory paths configured) then CIFS configuration will not be transitioned and only NFS configuration will be transitioned.

To transition the root volume of the vFiler with NFS only configuration transition, and ignoring the CIFS configuration transition, the default HOME share (for home-directory within the vFiler root volume) that is created when the CIFS server is configured should be deleted.

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 clustered Data 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 clustered Data 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 clustered Data 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 clustered Data 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.
For more information managing export policies and rules, see the *Clustered Data ONTAP File Access Management Guide for NFS*.

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 clustered Data 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 clustered Data 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: 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::>

```

For the NFS client 192.168.10.10 to access the qtree, the NFS client 192.168.10.10 must have read-only access to the qtree's parent volume.

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

```

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

```

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

```

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

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

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

                                Vserver: std_22
                                Policy Name: std_2226
                                Rule Index: 1
                                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 clustered Data 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 clustered Data 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.

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

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

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

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

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

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

Considerations for transitioning UNIX users and groups

The 7-Mode Transition Tool precheck fails if the total number of UNIX users and groups that are being transitioned and the UNIX users and groups existing on the SVM exceeds the limits defined in clustered Data ONTAP.

You can perform one of the following actions to resolve this error:

- Transition the volumes by excluding the transition of UNIX users and groups
You can use the `transition property-set -s subproject_name -p ignore-unix-users-groups-transition -v true` command to exclude the transition of UNIX users and groups.

- Move UNIX users and groups to NIS or LDAP servers before transition

Related information

Clustered Data ONTAP 8.3 File Access Management Guide for NFS

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.

- All the prerequisites for CIFS transition must be met.
Prerequisites for transitioning CIFS configurations on page 50
- 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 because 7-Mode Transition Tool does not support the transition of local users and groups to clustered Data ONTAP 8.2.

- 7-Mode Transition Tool supports the transition of local users and groups to clustered Data ONTAP 8.2.1.

However, 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 such transitions to proceed.

While transitioning local users and groups from a given 7-Mode controller or a vFiler unit to a given SVM, local users and groups from the first transition subproject are transitioned. In the subsequent transition of subprojects 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.

For more information, see the *Clustered Data ONTAP Release Notes*.

- 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 clustered Data ONTAP 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 using the following clustershell CLI command: `cifs users-and-groups local-user set-password -vserver <vserver name> -user-name <user-name>`.

- The 7-Mode Transition Tool blocks the transition of local users and groups to clustered Data ONTAP during the precheck phase, if a local group has multiple member system identifiers (SIDs) mapped to a single domain user or group on the 7-Mode system.

On encountering such an 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.

Space considerations when transitioning SAN volumes

You must ensure that there is sufficient space available in the source and destination 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.

Prior to cutover, you can use the `df -h` command on the 7-Mode volume to verify if a 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 transition of LUNs fails due to lack of space on the destination volume, the following EMS message is generated: `LUN.vol.proc.fail.no.space: Processing for LUNs in volume vol1 failed due to lack of space.`

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

If there are volumes containing space-reserved LUNs, growing the volume does provide sufficient space. In such cases, the amount of space that has to be added is the size of the Snapshot reserve for the volume. After space is added to the destination volume, you can use the `lun transition start` command to transition the LUNs.

Related information

[Data ONTAP 8.2 Data Protection Online Backup and Recovery Guide for 7-Mode](#)

Considerations for deduplication and compression

For 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 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 and 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. This avoids sending undeduplicated or uncompressed data and additional temporary metadata files over the network to the destination volume.

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

Starting 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 version 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 if 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 clone hierarchy and storage efficiency are lost in the process. Therefore, the clones are split from the parent volumes and are transitioned as FlexVol volumes to the destination cluster.

FlexClone volumes cannot be created from Snapshot copies that are transitioned from 7-Mode. You can create FlexClone volumes from new Snapshot copies that are created after the volume is transitioned to clustered Data ONTAP.

Considerations for transitioning 7-Mode volumes to Storage Virtual Machines (SVMs) in the nondefault IPspace

You must be aware of certain considerations when transitioning 7-Mode volumes to Storage Virtual Machines (SVMs) that belong to a nondefault IPspace.

- If the target SVM is in a nondefault IPspace, you must not use the 7-Mode Transition Tool for transitioning 7-Mode IP addresses or configuring new LIFs on the SVM.
- When creating or editing a transition project by using the 7-Mode Transition Tool, you must omit the interface mapping step, if the target SVM is in a nondefault IPspace.
You can use OnCommand System Manager or the Data ONTAP CLI to configure the LIFs on the SVM after the transition.
- If the target SVM is in a nondefault IPspace and 7-Mode IP addresses are selected for transition in the project, the precutover operation fails during transition.
You must remove the IP addresses from the project and run the precutover operation again. You can create the LIFs manually on the SVM after transition.

Considerations for Snapshot copies

After transition, a Snapshot copy can be restored by using SnapRestore (file and volume). However, restoring Snapshot copies is not supported if the Snapshot copy contains qtree SnapMirror or LUN data. Additionally, FlexClone volumes cannot be created from the Snapshot copies that are transitioned to clustered Data ONTAP.

Considerations for quotas

Quotas are not applied after the precutover phase of transition. They are applied only during the storage cutover phase.

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

- Read-only mode: Quotas are not applied in the precutover read-only mode on the clustered Data 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 clustered Data ONTAP system so that you can test them in clustered Data ONTAP. However, the quotas are removed during resynchronization (after testing is completed) of the clustered Data ONTAP volumes. The quotas are applied again during the storage cutover phase.

Considerations for transitioning LDAP and DNS configurations

You must be aware of how the DNS and LDAP configurations in 7-Mode are transitioned and applied in clustered Data ONTAP.

Considerations for LDAP transition

- If multiple base and scope values are set for the `ldap . base`, `ldap . base . passwd`, `ldap . base . group`, or `ldap . base . netgroup` options 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.
- If multiple scope values are set for the `ldap . base`, `ldap . base . passwd`, `ldap . base . group`, or `ldap . base . netgroup` options and if you are transitioning to clustered Data ONTAP 8.2.2, only one value for each option is transitioned.
- If separate base and scope values are specified for user mapping (`ldap . usermap . base`) and user password (`ldap . base . passwd`) lookups in the 7-Mode system, only the base and scope values for the user password are transitioned, and these values are used for both user mapping and user password lookups in clustered Data ONTAP. This can cause security issues. You should manually add the base and scope values for user mapping to the user DN option in clustered Data ONTAP after transition, if required.

Considerations for DNS transition

A maximum of six domain names and a maximum of three name servers for DNS configurations are supported in clustered Data ONTAP. If the number of domains or name servers in 7-Mode exceeds the limit, only the first six domains and first three nameservers are transitioned.

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 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.

About backing up volumes in transition to tape

You can use the dump engine to perform tape backup and restore operations on a clustered Data ONTAP volume that is in transition during all transition phases except when the baseline transfer (SnapMirror initialization) is in progress. Whenever the volume changes state from read/write to read-only or from read-only to read/write, you must perform a baseline backup.

You must perform a baseline backup on the transitioned volume in the following scenarios:

- After SnapMirror initialization, when a SnapMirror relationship is established between the 7-Mode volume and the clustered Data ONTAP volume
The clustered Data ONTAP volume is now in the read-only state.
- When the SnapMirror relationship between the 7-Mode volume and the clustered Data ONTAP volume is broken
The clustered Data ONTAP volume is now in the read/write state.

For a 7-Mode volume SnapMirror relationship, the additional scenarios in which you must perform a baseline backup on the clustered Data ONTAP secondary volume are as follows:

- When a SnapMirror relationship is established between the 7-Mode primary volume and the clustered Data ONTAP secondary volume
The clustered Data ONTAP secondary volume is in the read-only state.
- After the 7-Mode primary volume is transitioned to a clustered Data ONTAP volume, the SnapMirror relationship between the 7-Mode primary volume and the clustered Data ONTAP secondary volume is broken
The clustered Data ONTAP secondary volume is now in the read/write state.
- When a SnapMirror relationship is established between the clustered Data ONTAP primary and secondary volumes

For more information about the dump engine, see the *Clustered Data ONTAP Data Protection Tape Backup and Recovery Guide*.

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 were not transitioned, before enabling client access to the clustered Data 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.
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, Snapshot policies, export policies and volume efficiency settings, that might need to be consolidated or deleted manually.
3. SAN transition: manually reconfigure the hosts.
[7-Mode Transition Tool 2.1 SAN Host Transition and Remediation Supplemental Guide](#)
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 clustered Data ONTAP volumes.

Related information

[NetApp Documentation: Data ONTAP 8 \(current releases\)](#)

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 see which LUNs failed transition to clustered Data ONTAP, and then determine a corrective course of action.

Steps

1. Use the `lun transition 7-mode show` command to see which LUNs failed.
2. Examine the EMS logs and determine the corrective action that must be taken.
3. Take the steps necessary to correct the failure.
4. Use the `lun transition start` command to complete the transition.
5. Use the `lun show` command to verify that the failed LUN was successfully transitioned.

This command displays only the LUNs that were successfully transitioned to clustered Data ONTAP.

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.

Before you begin

Storage cutover must have completed successfully.

Step

1. For generating LUN mapping information, run the following command from the CLI:

```
transition export lunmap -s session_name -o file_path
```

Example

```
C:\Program Files\NetApp\7-Mode Transition Tool>transition 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
```

After you finish

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

[7-Mode Transition Tool 2.1 SAN Host Transition and Remediation Supplemental Guide](#)

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.

Note: 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 clustered Data 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 clustered Data ONTAP 8.3.1 and later

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 clustered Data ONTAP.
- Create a clone of the volume in clustered Data 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.

Related information

[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.


Before you begin

Storage cutover for the 7-Mode active site must be 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.

Example



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.

About this task

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.

Step

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 Snapshot policies and cron schedules from transitioned volumes

The 7-Mode Transition Tool generates unique schedules and Snapshot policies per transitioned volume. This can lead to a large number of cron schedules and Snapshot policies being created after the transition, which makes management more complicated than it needs to be.

Step

1. For better management, consolidate the cron schedules and Snapshot policies after transition by using the `job schedule cron` and `volume snapshot policy` commands.

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.

Collecting tool logs

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

Steps

1. Click **Migrate**, and then click **Logs** in the **Home** tab.
2. On the **Transition Logs** page, click **Collect Tool Logs**.
3. In the **Transition Logs** dialog box, click **Yes** to download the logs to your local system.

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

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 the 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 session:

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

The following table lists the log files that are not related to a particular session:

Log file path	Contains information about...
transition-gui.log	Entries of all the actions performed by using the web interface
default/audit.log	• All the parameters—such as HTTP or HTTPS port and log directory path—that are used by the tool every time the 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 session

Log file path	Contains information about...
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 session
default/default/ zapi-outbound.log	Output of all the Data ONTAP APIs that are executed by the 7-Mode Transition Tool and which do not belong to any session
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 session
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.

Tool installation fails with an incompatible OS error

Installation of the 7-Mode Transition Tool fails with the error message 7-Mode Transition Tool is not compatible with this version of Windows.

Workaround

1. Right-click the NetApp_7ModeTransitionTool_Setup.exe file, and select **Properties**.
2. In the Compatibility tab, click **Change settings for all users**.
3. Ensure that the **Run this program in compatibility mode for:** check box is not selected, and click **Apply**.

Installation or uninstallation of the 7-Mode Transition Tool is blocked

When you use the 7-Mode Transition Tool installer or the uninstaller to perform any operations (including install, repair, reinstall, upgrade, uninstall, the following error message is displayed: Another instance of NetApp 7-Mode Transition Tool Setup or Uninstall is running. Only one instance can run at a time.

Workaround

Before you begin, ensure that there are no open instances or dialog boxes of the 7-Mode Transition Tool. Otherwise, you might have to perform the steps again.

1. Open **Registry** by clicking **Start**.
2. In **Search programs and files**, enter **RegEdit**.
3. When Registry Editor is open, locate the HKEY_LOCAL_MACHINE registry hive.
4. Expand the HKEY_LOCAL_MACHINE hive by clicking the expand icon to the left of the folder icon.
5. Continue to expand the registry keys and subkeys until you locate HKEY_LOCAL_MACHINE\SOFTWARE\NetApp\7-Mode Transition Tool.

6. Click **7-Mode Transition Tool**.
7. Right-click `INSTALLATION_STATUS`, and select **Modify**.
8. Change the value from `INSTALLER_RUNNING` to `INSTALLED`, and click **OK**.

You should now be able to perform any operation using the 7-Mode Transition Tool installer or the uninstaller.

Reinstalling or upgrading the 7-Mode Transition Tool fails on Windows system

Message

Another instance of NetApp 7-Mode Transition Tool setup or uninstall is running. Only one instance can run at a time.

Cause

Some of the files in the 7-Mode Transition Tool installation folder are open.

Corrective action

1. Close all the 7-Mode Transition Tool installation files.
2. Open the registry entry for 7-Mode Transition Tool:
 - a. Verify the `INSTALLATION_STATUS` field.
 - b. If the value is `INSTALLER_IS_RUNNING`, change it to `INSTALLED`.
3. Retry reinstalling or upgrading the tool.

Windows Program Compatibility Assistant notification displayed while installing the tool

Sometimes while installing the tool, you receive the Windows Program Compatibility Assistant notification that the program might not have installed correctly. This notification does not necessarily indicate that the program did not install correctly.

You can ignore this notification by clicking **This program installed correctly**.

Incorrect Data ONTAP version displayed if the 7-Mode system is upgraded after project creation

Issue

The Data ONTAP version of the 7-Mode system that is shown in the Storage Systems tab is incorrect.

Cause

If the 7-Mode system is upgraded after the project was created, the new Data ONTAP version is not updated in the Storage Systems tab.

Corrective action

In the Storage Systems tab, click  and reenter the credentials of the 7-Mode storage system.

The correct Data ONTAP version is displayed for the 7-Mode storage system.

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 Data 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

Select the volume after some time.

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.

Cannot use 7-Mode Transition Tool for transitioning 7-Mode volumes to Storage Virtual Machines (SVMs) in nondefault IPspace

Issue

The precutover operation fails during transition if you use the 7-Mode Transition Tool to create new LIFs when the target SVM is in a nondefault IPspace.

Message

Failed to create default route for the following routing groups.

Corrective action

1. Remove the existing LIFs from the transition project by clicking the **Remove** button in the Networking tab of the 7-Mode Transition Tool.
2. Run the precutover operation again.

Copyright information

Copyright © 1994–2015 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

- 32-bit aggregates
 - expanding to the 64-bit format [54](#)
- 32-bit data
 - finding and removing, after expanding to 64-bit [55](#)
 - preparing to transition when 7-Mode system includes [52](#)
- 64-bit
 - finding and removing 32-bit data after expanding to [55](#)
- 64-bit aggregates
 - expanding 32-bit aggregates to [54](#)
 - preparing to expand 32-bit to [52](#)
- 7-Mode
 - preparing for assessment [22](#)
 - preparing for transition [47](#)
- 7-Mode configurations
 - applying [69](#), [71](#)
 - applying NAS configurations [69](#)
 - copying [69](#), [71](#)
 - LUNs [71](#)
 - testing [69](#)
 - verifying [69](#)
- 7-Mode environments
 - data migration to clustered Data ONTAP [7](#)
 - transitioning volumes and configuration to clustered Data ONTAP [7](#)
- 7-Mode Snapshot copies
 - deleting from transitioned volumes [111](#)
- 7-Mode systems
 - adding for a transition session [88](#)
- 7-Mode systems with 32-bit data
 - preparing to transition [52](#)
- 7-Mode Transition Tool
 - configuration file [20](#)
 - creating a session through the CLI [88](#)
 - generating assessment reports by adding systems to [24](#)
 - installing [13](#)
 - installing on Linux [18](#)
 - installing on Windows [15](#)
 - log files [113](#)
 - logging in using the web interface [17](#)
 - network port requirements [49](#)
 - system requirements [13](#)
 - uninstalling [13](#), [16](#)
 - uninstalling on Linux [19](#)
 - upgrading [13](#), [15](#)
 - web interface
 - logging in to 7-Mode Transition Tool [17](#)
- 7-Mode Transition Tool is not compatible with this
 - version of Windows
 - workaround [114](#)
- 7-Mode versions
 - supported for migration [43](#)
- 7-Mode volume transition
 - tool cannot transition to SVMs in nondefault IPspace [117](#)
- 7MTT

See 7-Mode Transition Tool

A

- adding systems for migration [62](#)
- aggregates
 - expanding to the 64-bit format [54](#)
 - finding and removing 32-bit data after expanding to 64-bit [55](#)
- applications
 - supported versions for transition assessment [22](#)
- assessment
 - host and storage inventory information [22](#)
 - how to use the executive summary for transition [26](#)
 - preparing for transition [22](#)
 - preparing the 7-Mode system [22](#)
 - when to perform while transitioning [7](#)
- assessment limits
 - maximum number of controllers, hosts [12](#)
- assessment reports
 - generating by adding systems to the 7-Mode Transition Tool [24](#)

B

- backup of volumes in transition
 - using dump [106](#)
- baseline data copy
 - data copy from 7-Mode
 - starting [68](#)
 - starting [68](#)
 - transition phase described [29](#)
- best practices
 - during data copy [79](#)
 - for data copy [79](#)
- bidirectional SnapMirror relationships
 - how to transition [87](#)
- browsers
 - supported versions for the web interface [13](#)

C

- checklist
 - for transition [58](#)
- CIFS
 - prerequisites to transition [50](#)
 - scenario for transition [87](#)
- CIFS migration
 - workflow [40](#)
- client access
 - post-transition tasks before enabling [108](#)
 - volumes in data protection relationship [34](#)
- clustered Data ONTAP
 - transitioning data and configuration from a 7-Mode environment [7](#)
- clusters
 - adding [62](#)
 - preparing for transition [45](#)

- collection
 - host and storage inventory information [22](#)
- commands
 - for managing transition service on Linux systems [20](#)
- commands to manage
 - credentials [94](#)
 - jobs [94](#)
 - LIFs [94](#)
 - session properties [94](#)
 - SnapMirror relationships [94](#)
 - transition operations [94](#)
 - transition sessions [94](#)
 - volume pairs [94](#)
- comments
 - how to send feedback about documentation [120](#)
- compatibility
 - verifying, for transition [91](#)
- compression
 - considerations for transition [104](#)
- configuration file
 - modifying [20](#)
- configuration transition phase
 - transition of stand-alone volumes [30](#)
- configurations
 - unsupported features and volumes for transition [8](#)
- considerations
 - for quotas while transitioning [105](#)
 - for transitioning 7-Mode volumes to clustered Data ONTAP in MetroCluster configurations [51](#)
 - for transitioning DNS and LDAP [106](#)
 - for transitioning FlexClone volumes [105](#)
 - space requirements for transitioning LUNs [104](#)
- considerations during 7-Mode volume transition
 - SVM in nondefault IPspace [105](#)
- considerations for NFS [100](#)
- considerations for NFS exports transition [100](#)
- considerations for NFS transition [100](#)
- considerations for transition
 - UNIX users and groups [102](#)
- controllers
 - adding [62](#)
 - assessing transition readiness of [23](#)
- creating
 - a transition session [88](#)
- credentials
 - commands to manage [94](#)
- cron schedules on transitioned volumes
 - consolidating to ease management [112](#)
- cutover
 - completing transition for [94](#)
 - guidelines for performing [99](#)
 - performing post-transition tasks after [108](#)
- cutover phase
 - volumes in a data protection relationship [34](#)

D

- dashboard
 - aborting subprojects [85](#)
 - prechecks [66](#)
- data
 - finding and removing 32-bit, after expanding to 64-bit [55](#)

- data copy
 - considerations for scheduling [79](#)
 - initiating [92](#)
 - priority for transfers [79](#)
- data copy phase
 - transition of stand-alone volumes [30](#)
 - volumes in a data protection relationship [34](#)
- data copy schedule
 - creating [74](#)
 - managing [74](#)
- data copy schedules
 - considerations for [98](#)
 - creating [88](#)
 - deleting [80](#)
 - editing [80](#)
- data copy update
 - scheduling [74](#)
- Data ONTAP
 - 7-Mode versions supported [43](#)
 - 7-Mode versions supported for transition assessment [22](#)
- data resynchronization operation
 - scheduling [74](#)
- deduplication
 - considerations for transition [104](#)
- deleting
 - SnapMirror transfer schedule [80](#)
- DNS
 - considerations for transition [106](#)
- documentation
 - how to receive automatic notification of changes to [120](#)
 - how to send feedback about [120](#)
- downloading
 - the log files [113](#)
- dump support for transition
 - understanding [106](#)

E

- editing subprojects
 - using web interface [74](#)
- error messages
 - 7-Mode Transition Tool installer is blocked [114](#)
- executive summary
 - how you can use it for transition assessment [26](#)
- expanding
 - aggregates to the 64-bit format [54](#)

F

- failures
 - recovering from LUN transition [108](#)
- FC transition
 - workflow [40](#)
- features
 - unsupported for transition [8](#)
- feedback
 - how to send comments about documentation [120](#)
- files
 - collecting the logs [113](#)
- FlexClone volumes

- considerations for transition [105](#)
- FlexVol volumes
 - finding and removing 32-bit, after expanding aggregates to 64-bit [55](#)
- flowcharts
 - CIFS migration [40](#)
 - NFS migration [39](#)
- functionality
 - unsupported 7-Mode SAN, in clustered Data ONTAP [10](#)

G

- giveback
 - impact on transition [106](#)
- guidelines
 - for creating data copy schedules [79](#)

H

- hosts
 - assessing transition readiness of [23](#)
- hosts, supported
 - versions for transition assessment [22](#)

I

- importing
 - inventory report XML file for transition assessment [26](#)
- incompatible OS error
 - tool installation fails [114](#)
- incorrect 7-Mode version displayed after project creation
 - troubleshooting [115](#)
- information
 - how to send feedback about improving documentation [120](#)
- installation
 - workaround for error message during [114](#)
- installation fails
 - incompatible OS error [114](#)
- installing
 - 7-Mode Transition Tool [13](#)
 - 7-Mode Transition Tool, Linux [18](#)
 - 7-Mode Transition Tool, Windows [15](#)
- inventory information
 - collecting and assessing [22](#)
- inventory reports
 - importing for transition assessment [26](#)
- IP addresses
 - adding for a transition session [88](#)
- iSCSI transition
 - workflow [41](#)

J

- jobs
 - commands to manage [94](#)

L

- LDAP
 - considerations for transition [106](#)
- licenses
 - requirements for transition [43](#)
- LIF
 - See* logical interfaces
- LIFs
 - commands to manage [94](#)
- limitations
 - SnapDrive-managed 7-Mode Snapshot copies [110](#)
 - SnapManager-managed 7-Mode Snapshot copies [110](#)
- limits
 - number of controllers and hosts for assessment [12](#)
- Linux systems
 - commands for managing transition service [20](#)
- log files
 - 7-Mode Transition Tool [113](#)
- logging in
 - to the 7-Mode Transition Tool [17](#)
- logical interfaces
 - managing [83](#)
 - See also* LIFs
- login access
 - for users not member of Administrator group [16](#)
- logs
 - collecting 7-Mode transition Tool logs [113](#)
- LUN.vol.proc.fail.no.space
 - cause and corrective action if transition fails [104](#)
- LUNs
 - recovering from transition failures [108](#)

M

- MetroCluster configuration
 - transitioning after switchback or during switchover [116](#)
- MetroCluster configurations
 - requirements and considerations for transitioning 7-Mode volumes to clustered Data ONTAP [51](#)
- MetroCluster transition
 - post-transition task [111](#)
- migration
 - Data ONTAP version requirements [43](#)
 - downloading the log files [113](#)
 - preparing for [43](#)
 - when to perform while transitioning [7](#)
 - workflow for CIFS [40](#)
 - workflow for NFS [39](#)
- migration phases
 - baseline data copy [29](#)
 - precutover [29](#)
 - preparation [29](#)
 - storage cutover [29](#)
- mixed-format volumes
 - finding and removing after expanding aggregates to 64-bit [55](#)
- modifying
 - SnapMirror transfer schedule [80](#)
- multiple paths
 - considerations for data copy using [98](#)

N

network ports
requirements [49](#)

NFS
prerequisites to transition [49](#)

NFS migration
workflow [39](#)

NFS only transition
NFS only configuration transition [99](#)

O

on-demand SnapMirror updates
performing [81](#)
operational changes
when to perform while transitioning [7](#)
overview
transition [7](#)

P

pausing
transition subprojects [84](#)
phases
stand-alone volumes, transition [30](#)
volumes in data protection relationship, transition [34](#)
planning
when to perform while transitioning [7](#)
ports
requirements [49](#)
post transition steps
transition of stand-alone volumes [30](#)
post-transition
creating cron job schedules [111](#)
post-transition limitation
7-Mode Snapshot copies of LUNs [110](#)
post-transition tasks
for configuring features not transitioned [108](#)
prechecks
error [67](#)
information [67](#)
notification [67](#)
running [66](#), [91](#)
severity levels [67](#)
warning [67](#)
precutover
transition phase described [29](#)
precutover phase
volumes in a data protection relationship [34](#)
preparation
for transition, checklist [58](#)
transition phase described [29](#)
preparation phase
transition of stand-alone volumes [30](#)
volumes in a data protection relationship [34](#)
preparing
7-Mode system for assessment [22](#)
7-Mode system for transition [47](#)
cluster for transition [45](#)
for migration [43](#)
protocols

preparing the 7-Mode system for assessment [22](#)
provisioning SVM
using System Manager [52](#)

Q

quotas
considerations for transition [105](#)

R

reinstallation fails
on Windows, corrective action [115](#)
removing
volume from subproject [83](#)
reports
importing inventory reports for transition assessment [26](#)
reports, assessment
generating by adding systems to the 7-Mode Transition Tool [24](#)
requirements
7-Mode system and hosts supported for transition assessment [22](#)
7-Mode Transition Tool [49](#)
for transitioning 7-Mode volumes to clustered Data ONTAP in MetroCluster configurations [51](#)
for transitioning CIFS configurations [50](#)
for transitioning NFS configurations [49](#)
licenses for transition [43](#)
requirements for transition
transition service [44](#)
requirements, system
for the 7-Mode Transition Tool [13](#)
restrictions
on transition [44](#)
resuming
transition subprojects [84](#)

S

SAN
post-transition task [109](#)
post-transition tasks [108](#)
unsupported 7-Mode features in clustered Data ONTAP [10](#)
unsupported features and volumes for transition [8](#)
SAN transition
LUN mapping [109](#)
space considerations [104](#)
SAN volumes
space requirements for transitioning LUNs [104](#)
schedules
creating for data copy [88](#)
guidelines for creating data copy [79](#)
secondary subprojects
aborting [85](#)
sessions
scenarios for creating [87](#)
SnapMirror baseline copy
scheduling [74](#)
SnapMirror relationships

- adding information for a transition session [88](#)
 - commands to manage [94](#)
 - considerations for transition [98](#)
 - scenario for transition [87](#)
- SnapMirror schedule
 - creating [74](#)
- SnapMirror schedules
 - managing [74](#)
- SnapMirror transfer schedules
 - deleting [80](#)
 - editing [80](#)
- SnapMirror updates
 - performing on-demand [81](#)
- Snapshot copies
 - 7-Mode, post-transition limitation [110](#)
 - finding and removing 32-bit, after expanding aggregates to 64-bit [55](#)
- Snapshot policies on transitioned volumes
 - consolidating to ease management [112](#)
- space considerations
 - when transitioning SAN volumes [104](#)
- stand-alone subprojects
 - aborting [85](#)
- storage cutover
 - transition phase described [29](#)
- storage efficiency
 - considerations for transition [104](#)
- subprojects
 - aborting [85](#)
 - adding volumes [74](#)
 - editing [74](#)
 - managing [74](#)
 - modifying transition configuration [74](#)
 - pausing and resuming [84](#)
 - removing volumes [74](#)
- suggestions
 - how to send feedback about documentation [120](#)
- switchback
 - transition a MetroCluster configuration [116](#)
- switchover
 - transition a MetroCluster configuration [116](#)
- system requirements
 - for the 7-Mode Transition Tool [13](#)

T

- takeover
 - impact on transition [106](#)
- throttle limit
 - specifying [74](#)
- throttle limits
 - specifying [81](#)
- transfers
 - priority for [79](#)
 - transition considerations for concurrent SnapMirror [98](#)
- transition
 - assessing controllers and hosts are ready for [23](#)
 - assessing hosts and controllers for [22](#)
 - best practices during the different phases of [29](#)
 - checklist [58](#)
 - CIFS configuration prerequisites [50](#)

- commands to configure transition service on Linux systems [20](#)
 - commands to manage [94](#)
 - configuration file for [20](#)
 - considerations for deduplication and compression [104](#)
 - considerations for FlexClone volumes [105](#)
 - considerations for local users and groups [103](#)
 - considerations for quotas [105](#)
 - considerations for SnapMirror [98](#)
 - downloading the log files [113](#)
 - examples for creating sessions [87](#)
 - generating assessment reports by adding systems to the 7-Mode Transition Tool [24](#)
 - guidelines for cutover [99](#)
 - how to use the execute summary for assessment [26](#)
 - impact of takeover and giveback [106](#)
 - importing inventory report XML file for readiness check [26](#)
 - incorrect Data ONTAP version displayed, troubleshooting [115](#)
 - license requirements [43](#)
 - local users and groups
 - considerations during CIFS transition [103](#)
 - NFS configuration prerequisites [49](#)
 - overview [7](#)
 - performing manual tasks after [108](#)
 - phases of [34](#)
 - prechecks [67](#)
 - preparing the 7-Mode system [47](#)
 - preparing the cluster for [45](#)
 - project definition [11](#)
 - removing volumes from a subproject [83](#)
 - restrictions on [44](#)
 - session [87](#)
 - session definition [11](#)
 - stand-alone volumes [30](#)
 - subproject definition [11](#)
 - terminology [11](#)
 - unsupported features and volumes for [8](#)
 - unsupported SAN features in clustered Data ONTAP [10](#)
 - using the 7-Mode Transition Tool (CLI) [87](#)
 - volumes in a data protection relationship [34](#)
- transition assessment
 - 7-Mode system and hosts supported for [22](#)
- transition complete phase
 - transition of stand-alone volumes [30](#)
- transition considerations
 - DNS and LDAP [106](#)
 - Snapshot copies [105](#)
- transition failures
 - recovering from LUN [108](#)
- transition issues
 - cannot select a secondary volume [116](#)
- transition limits
 - number of active projects [12](#)
 - number of volumes in a subproject [12](#)
- transition operations
 - commands to manage [94](#)
- transition project
 - completing [72](#)
 - creating [63](#)

- deleting [86](#)
 - managing [74](#)
 - running [57](#)
 - workflow [57](#)
- transition projects
 - prechecks [66](#)
- transition service
 - managing on Linux systems [20](#)
- transition session
 - applying configurations [93](#)
- transition sessions
 - commands to manage [94](#)
 - completing [94](#)
 - creating [88](#)
 - initiating data copy [92](#)
 - managing [88](#)
 - starting [92](#)
 - verifying compatibility [91](#)
- transition subprojects
 - completing [72](#)
 - pausing [84](#)
 - resuming [84](#)
- transition-tool.conf
 - modifying [20](#)
- transitioned volumes
 - deleting 7-Mode Snapshot copies [111](#)
- transitioning
 - from 7-Mode environments to clustered Data ONTAP [7](#)
 - MetroCluster configuration that failed [116](#)
 - preparing to transition [52](#)
- transitioning 7-Mode volumes to SVMs in nondefault IPspace
 - cannot use 7-Mode Transition Tool [117](#)
- transitioning volumes
 - excluding CIFS local users and groups [82](#)
 - excluding configurations [82](#)
 - netgroup configurations [82](#)
 - UNIX user and groups [82](#)
- transitions
 - requirements and considerations for 7-Mode to clustered Data ONTAP volumes in MetroCluster configurations [51](#)
- troubleshooting
 - incorrect Data ONTAP version displayed if the 7-Mode system is upgraded [115](#)
- troubleshooting issues
 - reinstallation or upgrade fails [115](#)
- twitter
 - how to receive automatic notification of documentation changes [120](#)

U

- uninstalling
 - 7-Mode Transition Tool [16](#)

- 7-Mode Transition Tool, Linux [19](#)
- UNIX users and groups
 - considerations for transition [102](#)
- unsupported features
 - for transition [8](#)
- updates
 - performing on-demand SnapMirror [81](#)
- upgrade
 - deciding if required before transition [52](#)
- upgrade fails
 - on Windows, corrective action [115](#)
- upgrading
 - 7-Mode Transition Tool [15](#)
- users
 - enabling login for users not in Administrator group [16](#)

V

- volume mapping
 - no aggregates assigned [117](#)
- volume pairs
 - commands to manage [94](#)
- volume transitions
 - requirements and considerations for 7-Mode to clustered Data ONTAP, in MetroCluster configurations [51](#)
- volumes
 - adding [63](#)
 - adding for a transition session [88](#)
 - finding and removing 32-bit and mixed-format, after expanding aggregates to 64-bit [55](#)
 - LIFs
 - configuring [63](#)
 - logical interface
 - configuring [63](#)
 - removing from a subproject during transition [83](#)
 - scenarios for selecting them for a session [87](#)
 - space requirements for transitioning LUNs [104](#)
 - unsupported for transition [8](#)
- volumes in transition
 - backing up [106](#)
- Vservers
 - See SVMs

W

- web interface
 - troubleshooting [113](#)
- workflow
 - FC transition [40](#)
 - iSCSI transition [41](#)
- workflows
 - CIFS migration [40](#)
 - NFS migration [39](#)