



7-Mode Transition Tool 1.4

Data and Configuration Transition Guide

For Transitioning to Clustered Data ONTAP®



Contents

Transition overview	6
What the 7-Mode Transition Tool does	6
Features and volumes not supported for transition	7
Limits for transition	9
Installing, upgrading, and uninstalling the 7-Mode Transition Tool	11
System requirements for 7-Mode Transition Tool	11
Installing the 7-Mode Transition Tool	12
Modifying the configuration options of the 7-Mode Transition Tool	13
Logging in to the tool	14
Upgrading the 7-Mode Transition Tool	15
Uninstalling the 7-Mode Transition Tool	16
Collecting and assessing the inventory information	17
Storage and host version requirements for transition assessment	17
Preparing the 7-Mode systems and hosts for transition assessment	19
Assessing controllers and hosts	20
Generating an assessment report by adding systems to the 7-Mode Transition Tool	21
Generating an assessment report by importing the inventory report XML ...	23
How you can use the assessment executive summary for transition assessment	24
Migration process and phases	26
How you transition a stand-alone volume	28
How you transition volumes in a SnapMirror relationship	31
NFS migration workflow	36
CIFS migration workflow	37
Preparing for migration	39
Data ONTAP version requirements for migration	39
License requirements for transition	39
7-Mode Transition Tool service must be running	40
Restrictions for transition	40
Preparing the cluster for transition	41
Preparing the 7-Mode system for transition	44
Port requirements for communicating with the 7-Mode Transition Tool	46

Prerequisites for transitioning NFS configurations	47
Prerequisites for transitioning CIFS configurations	48
Migrating data and configuration from 7-Mode volumes	50
Transition preparation checklist	51
Adding controllers and clusters	56
Creating a transition project	57
Running prechecks	61
Severity levels for precheck messages	62
Starting baseline data copy	63
Applying 7-Mode configurations	65
Completing a transition project	67
Managing a transition project	69
Editing a subproject	69
Managing SnapMirror transfers and schedule	69
Creating a data copy schedule for SnapMirror transfers	70
Considerations for creating a data copy schedule	74
Editing or deleting a data copy schedule for SnapMirror transfers	75
Performing on-demand SnapMirror updates	77
Transitioning volumes by excluding certain 7-Mode configurations	78
Managing logical interfaces	79
Removing volumes from a subproject	80
Pausing and resuming a subproject	81
Aborting a subproject	81
Deleting a project	82
Considerations for transition	83
Considerations for using SnapMirror for transition	83
Guidelines for deciding when to perform cutover	85
NFS only configuration transition	85
How NFS exports are transitioned	85
Considerations for transitioning CIFS local users and groups	89
Considerations for deduplication and compression	90
Considerations for FlexClone volumes	91
Considerations for Snapshot copies	91
Considerations for quotas	91
Considerations for transitioning thin-provisioned volumes	91
Considerations for transitioning LDAP and DNS configurations	92

Impact of takeover and giveback on transition	92
About backing up volumes in transition to tape	93
NFS configuration settings that are transitioned	93
CIFS configuration settings that are transitioned	95
Troubleshooting issues	98
Collecting tool logs	98
Log files for the 7-Mode Transition Tool	98
Cannot select a secondary volume from the Volume selection pane	99
Cannot select a volume for transition if the tool fails to retrieve the volume information	100
Cannot proceed from the Volume mapping pane if the selected SVM has no aggregate	100
Windows Program Compatibility Assistant notification displayed while installing the tool	100
Installation or uninstallation of the 7-Mode Transition Tool is blocked	100
Volume goes to an inconsistent state after transition is paused	101
Commands for managing transition	103
Transition terminology	107
Copyright information	109
Trademark information	110
How to send your comments	111
Index	112

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 does

The 7-Mode Transition Tool enables you to assess 7-Mode controllers, hosts, and applications for transition and perform copy-based transitions of Data ONTAP 7G and 7-Mode FlexVol volumes and configurations to new hardware that is running clustered Data ONTAP 8.2, 8.2.1, or 8.2.2, with minimum client disruption and retention of storage efficiency options.

Attention: Using the 7-Mode Transition Tool, you can transition only network-attached storage (NAS) environments to clustered Data ONTAP.

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

7-Mode Transition Tool works at the volume-level. 7-Mode Transition Tool uses SnapMirror technology for transferring data. It sets 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 only supported for

the purpose and duration of a migration that is part of a transition to clustered Data ONTAP. It is not possible to resynchronize data from a clustered Data ONTAP volume back to a 7-Mode volume.

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 any differences in the clustered Data ONTAP version selected for transition
- Migrate data and configuration from systems running different versions of Data ONTAP operating in 7-Mode to clustered Data ONTAP versions 8.2, 8.2.1, and 8.2.2
- Transition existing standalone 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 in clustered Data ONTAP before disconnecting client access
This allows for verifying the configurations in clustered Data ONTAP and minimizing the downtime.

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, such as synchronous SnapMirror relationships, because either some features might not be available in clustered Data ONTAP or the tool cannot transition these features.

Note: You can transition a 7-Mode volume only to an SVM with FlexVol volumes, and not to an SVM with Infinite Volume.

You cannot transition the following 7-Mode volumes:

- Volumes with LUNs and LUN clones (SAN transition)
- Restricted or offline volumes
- Traditional volumes
- Volumes with NFS-to-CIFS character mapping (charmap)
- Volumes with Storage-Level Access Guard configurations
- SnapLock volumes

8 | Data and Configuration Transition Guide

- FlexCache volumes
- 7-Mode volumes in a 64-bit aggregate to a 32-bit aggregate in clustered Data ONTAP
- 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:

- IPv6 configurations
- SnapVault relationships
While SnapVault source volumes can be migrated, you must first break the SnapVault relationship. SnapVault relationships or a volume that is the destination of a SnapVault relationship cannot be migrated.
- FPolicy configurations
- BranchCache configurations
- Antivirus configurations
- FCoE 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
- Qtree SnapMirror relationships
- MetroCluster configurations
- Disaster recovery vFiler unit (vFiler DR)
Only the volumes in the source vFiler unit can be transitioned.
- 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 40

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	8 (each project can have either a stand-alone subproject or a primary and a secondary subproject) Note: If you have eight active projects, you can replicate data from eight 7-Mode storage systems and their associated SnapMirror destinations.

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

You can download and install the 7-Mode Transition Tool on a system that is running a 64-bit Windows 7 operating system. You can also upgrade the tool from the 1.3 version to the 1.4 version. When required, you can uninstall the 7-Mode Transition Tool.

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

Software requirements

- 64-bit Windows 7 Enterprise, Windows Server 2008 R2 Enterprise with SP1 onwards, or Windows Server 2012 R2 Standard
- 64-bit Oracle Java Runtime Environment (JRE) 64-bit JRE 1.7 update 55

Note: If the Windows host does not have JRE or has 32-bit JRE installed, the installer for the 7-Mode Transition Tool automatically installs 64-bit JRE 1.7 update 55. If the host has an earlier version of 64-bit JRE, the installer automatically updates it to JRE 1.7 update 55.

Server (web interface) 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 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
- 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.

Related concepts

[Port requirements for communicating with the 7-Mode Transition Tool](#) on page 46

Related information

[Interoperability Matrix: \[mysupport.netapp.com/matrix\]\(http://mysupport.netapp.com/matrix\)](http://mysupport.netapp.com/matrix)

Installing the 7-Mode Transition Tool

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, it is necessary to clear 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 39

[License requirements for transition](#) on page 39

[Port requirements for communicating with the 7-Mode Transition Tool](#) on page 46

Related information

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

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 Windows 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.
 - 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**.

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

Logging in to the tool

To transition volumes using the 7-Mode Transition Tool (web interface), you must log in to the application using your Windows credentials for the server on which the tool is installed.

About this task

It is a best practice to use a local user account (local to the system where 7-Mode Transition Tool is installed) with administrative privileges when logging in to the 7-Mode Transition Tool.

Steps

1. Navigate to the login page:

If you are accessing the web application...	Then...
From the server on which it was installed	Double-click the 7-Mode Transition Tool icon on your desktop.
From another computer	Navigate to <code>https://IP_address:port/transition</code> or <code>http://IP_address:port/transition</code> by using one of the supported web browsers. In this URL, <i>IP_address</i> is the IP address of the server and the <i>port</i> could 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.

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

Note: 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 using your Windows credentials for the server.

Upgrading the 7-Mode Transition Tool

You can upgrade to this release of the 7-Mode Transition Tool on your Windows system without uninstalling the previous version. Upgrading allows 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.
- The projects must be in completed or in an 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 1.3, you can upgrade to 7-Mode Transition Tool 1.4.

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 Site: mysupport.netapp.com](http://mysupport.netapp.com)

Uninstalling the 7-Mode Transition Tool

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.

Collecting and assessing the inventory information

You can collect inventory information from controllers, hosts, and host applications, 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, hosts attached to the storage systems, and host applications that are supported for transition assessment.

7-Mode versions for transition assessment

- Data ONTAP 7.3.3
- Data ONTAP 7.3.4
- Data ONTAP 7.3.5
- Data ONTAP 7.3.6
- Data ONTAP 7.3.7
- Data ONTAP 8.0
- Data ONTAP 8.0.1
- Data ONTAP 8.0.2
- Data ONTAP 8.0.3
- Data ONTAP 8.0.4
- Data ONTAP 8.0.5
- Data ONTAP 8.1
- Data ONTAP 8.1.2
- Data ONTAP 8.1.3

- Data ONTAP 8.1.4
- Data ONTAP 8.2
- Data ONTAP 8.2.1

Host versions

- Microsoft Windows Server 2008 R2 Datacenter
- Microsoft Windows Server 2008 R2 Enterprise
- Microsoft Windows Server 2012 R2 Datacenter
- Microsoft Windows Server 2012 R2 Standard
- Red Hat Enterprise Linux Server 6.5
- Red Hat Enterprise Linux Server 5.9
- VMware ESXi 5.1

Host clustering software

- Microsoft Cluster Server (MSCS)
- Red Hat Enterprise Linux Cluster

Host application versions

- Microsoft Exchange Server 2007
- SAP NetWeaver 7.3
- Oracle Database 10g
- VMware Workstation 9.0
- XenServer 6.1
- SnapDrive 6.5

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 HTTP on the 7-Mode system by using the following command:
options httpd.enable on
2. Enable SSL on the 7-Mode system:
 - a. Set up and start SSL by entering the following command:
secureadmin setup ssl
 - b. Enable SSL by entering the following command:
options ssl.enable on
 - c. If you are running Data ONTAP 8.x, enable SSLv3 by entering the following command:
options ssl.v3.enable on
3. 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.

4. Prepare your Windows host systems for the following:
 - Enable WMI access.
For more information about enabling WMI access, see the host documentation.

- Verify that you have installed the Microsoft Fibre Channel Information Tool (fcinfo) package and run the tool once on your Windows host system.
This tool enables you to collect the HBA configuration information of the host.
- If the system on which the 7-Mode Transition Tool is run does not belong to a domain, verify the following:
 - The host system belongs to a domain.
 - If the host has a local user, then the user name is in the following format:

```
SystemName\Username
```

5. Enable SSH on the Linux or ESXi host.

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

6. 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.

7. 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](https://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 21
- [Generating an assessment report by importing the inventory report XML](#) on page 23

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

You can collect inventory information for controllers and hosts by adding them to the tool and assess the features and functionalities of these systems and identify how these work in the clustered Data ONTAP version selected for transition by creating an assessment report.

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 which contains the system details, one system per line.

Each system details must be in the following format:

```
system type://user_name[:password]@host_name/IP
```

You can also provide the password later when performing the assessment task.

- 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 need to added to the tool again.

Steps

1. Log in to 7-Mode Transition Tool, and in the Collect and Assess section, click **Get Started**.
2. Click **Add Systems**, and 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. 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 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 AssessmentWorkbook and AssessmentExecutiveSummary reports are generated in XML format.

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

For viewing the AssessmentWorkbook in Microsoft Excel, you might have to enable macros in Excel.

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 98

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 98

How you can use the assessment executive summary for transition assessment

The transition executive summary provides a summary of the 7-Mode controllers, hosts, and host applications in your environment. It provides an assessment report of the current features and functionality that are used and helps in identifying the recommended transition tool 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. Based on the precheck results, the summary provides a recommendation about the best transition methodology for your workload.

If there are errors reported in the precheck, the report shows the affected controllers in red; these controllers cannot be transitioned to clustered Data ONTAP. For example, you cannot transition 7-Mode SnapLock volumes because SnapLock is not supported in clustered Data ONTAP.

You can also view the number of warnings generated for the controllers and volumes. You should review these warnings and resolve any issues before transitioning. Details of the transition precheck summary 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 the recommended transition method for each of the assessed controller. For most of the NAS configurations, 7-Mode Transition Tool is the recommended tool for transition. If you have SAN volumes, you should use a SAN transition tool, such as DTA2800.

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.

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 by 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 able to support the transition—for example, verifying that volumes are online and that there is network access between the systems. 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.

Although the precheck step and assessment steps done during the assessment appear 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 will be 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 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 for minimizing the time clients are disconnected.

Storage cutover

At a high-level, clients are disconnected, a final data transfer is performed, the SnapMirror relationship is broken, and then 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 command `cifs terminate -v user01` 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 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 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, baseline data copy, configuration transition, and transition complete. 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

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 warning reported by 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

This is the configuration transition phase. There are two modes for the precutover phase: precutover read-only and precutover read/write.

Precutover read-only mode

In this phase, configurations from the 7-Mode volumes are copied to the clustered Data ONTAP volumes. The 7-Mode Transition Tool performs the following tasks in this phase:

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

Precutover read/write mode

In this phase, configurations from the 7-Mode volumes are copied to the clustered Data ONTAP volumes. In the read/write mode, you can test your configurations in clustered Data ONTAP before storage cutover, thereby reducing downtime. This is an optional step.

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 in clustered Data ONTAP
6. After completing the configurations and data-access operations, resynchronizing the SnapMirror relationship between 7-Mode volumes and clustered Data ONTAP volumes

Storage cutover phase

1. Optional: Performing an on-demand SnapMirror update to reduce the down time 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. 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
6. Making 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 tasks:

1. Manual: Perform any steps to configure features that were not transitioned, as listed in the precheck report.
For example, IPv6 and FPolicy must be configured manually after transition.

2. 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.
3. Redirect client access to the clustered Data ONTAP volumes.

Related tasks

[Migrating data and configuration from 7-Mode volumes](#) on page 50

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 done 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, collects information about the 7-Mode primary system
 - For transitioning a primary volume, collects 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
6. Creating the data copy schedules for baseline and incremental transfers

You should not modify the objects (volumes, IP addresses, system information, and so on) on the controller after the subproject is prepared.

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

This is the configuration transition phase. There are two modes for the precutover phase: precutover read-only and precutover read/write.

Precutover read-only mode

In this phase, configurations from the 7-Mode volumes are copied to the clustered Data ONTAP volumes. The 7-Mode Transition Tool performs the following tasks in this phase:

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

Precutover read/write mode


In this phase, configurations from the 7-Mode volumes are copied to the clustered Data ONTAP volumes. In the read/write mode, you can test your configurations in clustered Data ONTAP before storage cutover, thereby reducing downtime. This is an optional step.

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 in clustered Data ONTAP
6. Resynchronizing the SnapMirror relationship between 7-Mode volumes and clustered Data ONTAP volumes based on the data copy schedules

Storage cutover (secondary volumes) 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. Establishing a SnapMirror relationship between the volumes on the 7-Mode primary system and the clustered Data ONTAP secondary volumes
 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.
6. 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
7. Taking the 7-Mode volumes offline


The following illustration depicts the transition of a secondary volume:



Storage cutover (primary volumes) 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. Breaking and deleting the SnapMirror relationship between the 7-Mode primary volume and the clustered Data ONTAP secondary volume
6. Setting up cluster peer and SVM peer relationships between the primary and secondary clusters
7. Setting up a SnapMirror relationship between the primary and secondary clustered Data ONTAP volumes
8. Resynchronizing the SnapMirror relationship between the clustered Data ONTAP volumes
9. 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
10. 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. 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.
3. Redirect client access to the clustered Data ONTAP volumes.

Related tasks

Migrating data and configuration from 7-Mode volumes on page 50

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 50

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 50

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 systems must be running the supported Data ONTAP versions.
- SnapMirror must be licensed on the 7-Mode system.
- Required feature licenses, if exists on the 7-Mode system, must be installed on the cluster.
- NTP server must be configured and 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 51

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, 8.2.1, or 8.2.2.

You can migrate volumes from systems running Data ONTAP 7.3.3 and later. For the updated list of Data ONTAP versions supported for migration by the 7-Mode Transition Tool, see the Interoperability Matrix.

Related information

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

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 machine 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.

- The tool does not allow you select an SVM that does not have NFS or CIFS enabled.
- 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.

- 7-Mode Snapshot copies that have LUNs cannot be used to restore a clustered Data ONTAP volume.
- 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.
- 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 will cause 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 the 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.
- An SFO policy must be enabled for the aggregates in the SVM to which the volumes will be transitioned.
- 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, 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, use the following command to verify that the cluster is reachable by using the cluster-management LIF:
ssh username@cluster_mgmt_IP
2. Use the `system services web show` command to verify that SSL is enabled on the cluster.

Example

```
cluster1::> system services web show
External Web Services: true
Status: online
HTTP Protocol Port: 80
HTTPS Protocol Port: 443
TLSv1 Enabled: true
SSLv3 Enabled: true
SSLv2 Enabled: true
```

3. Verify that HTTPS is allowed on the cluster-management LIF by performing the following steps:
 - 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 by performing the following steps:
 - 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. 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. 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

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

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 by using the following command:


```
license add xxxxxxxx
```

xxxxxxx is the license code you purchased.
 - b. Enable SnapMirror by entering the following command:

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 SSL on the 7-Mode system by performing the following steps:
 - a. Set up and start SSL by entering the following command:


```
secureadmin setup ssl
```
 - b. Enable SSL by entering the following command:


```
options ssl.enable on
```
 - c. If you are running Data ONTAP 8.x, enable SSLv3 by entering the following command:


```
options ssl.v3.enable on
```
4. Depending on the Data ONTAP version of your 7-Mode system, perform the following steps:
 - a. Allow SnapMirror traffic on all the interfaces by entering the following command:


```
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 by entering the following command:


```
options interface.blocked.mgmt_data_traffic off
```
5. 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 by using the following commands:


```
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 condition 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 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 needs to be transitioned as part of the sub-project 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.

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

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

Related concepts

[*NFS migration workflow*](#) on page 36

[*How NFS exports are transitioned*](#) on page 85

[*NFS only configuration transition*](#) on page 85

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 (Creating the CIFS server on the SVM by reusing the same CIFS server name of the 7-Mode system or vFiler unit.)	◦ The number of volumes in the 7-Mode system or vFiler unit to be transitioned must be less than or equal to 160 so that all the 7-Mode volumes can be transitioned in a single subproject. ◦ Before precutover, the CIFS server must be re-created on the 7-Mode system with a different name. ◦ During precutover, there must not be any active CIFS sessions on the 7-Mode system. ◦ During precutover, the new CIFS server must be up and running on the 7-Mode system. ◦ You can plan to perform the storage cutover immediately after precutover, to enable client access to the clustered Data ONTAP volumes.

If you want to..	Conditions
Use a new identity	◦ The CIFS server must be up and running on the 7-Mode system. ◦ CIFS must be running.

Related concepts

CIFS migration workflow on page 37

Considerations for transitioning CIFS local users and groups on page 89

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 the subproject is prepared.
- 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.

If you use the **Refresh** browser button, you are directed to the Login page.

The following image illustrates the migration process:



Steps

1. [Transition preparation checklist](#) on page 51
2. [Adding controllers and clusters](#) on page 56
3. [Creating a transition project](#) on page 57
4. [Running prechecks](#) on page 61
5. [Starting baseline data copy](#) on page 63
6. [Applying 7-Mode configurations](#) on page 65
7. [Completing a transition project](#) on page 67

Related concepts

[NFS migration workflow](#) on page 36

[CIFS migration workflow](#) on page 37

[How you transition a stand-alone volume](#) on page 28

[How you transition volumes in a SnapMirror relationship](#) on page 31

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 51
- [Licensing requirements](#) on page 52
- [SnapMirror requirements on the 7-Mode system](#) on page 52
- [Volume options on the 7-Mode system](#) on page 53
- [Managing access to the cluster](#) on page 53
- [Managing access to the 7-Mode system](#) on page 53
- [Networking requirements](#) on page 53
- [Port requirements](#) on page 54
- [NFS requirements](#) on page 55
- [CIFS requirements](#) on page 55

Data ONTAP version requirements

Item	Yes
Supported 7-Mode version NetApp Interoperability Matrix Tool	

Item	Yes
Your cluster must be running one of the following Data ONTAP versions: • Data ONTAP 8.2 • Data ONTAP 8.2.1 • Data ONTAP 8.2.2	

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	
Other feature licenses, if available on the 7-Mode system, are added to the cluster	

SnapMirror requirements on the 7-Mode system

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

Volume settings on the 7-Mode system

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

Managing access to the cluster

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

Managing access to the 7-Mode system

Item	Yes
SSL is enabled <code>secureadmin setup ssl</code> <code>options ssl.v3.enable on</code>	
HTTP is enabled <code>options httpd.admin.enable on</code>	
HTTPS is enabled <code>options httpd.admin.ssl.enable on</code>	

Networking requirements

Item	Yes
Cluster is reachable using the cluster-management LIF	

Item	Yes
One or more intercluster LIFs are set up on each node of the cluster For multipathing, two intercluster LIFs are required on each node.	
Static routes are created for the intercluster LIFs	
7-Mode system and cluster are reachable from the Windows system on which 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 added to the cluster	
For transitioning an AD-based Kerberos server, DNS entry exists for AD domain	
NFS added to the list of allowed protocols for the SVM.	
Clock skews between KDC and the cluster less than or equal to 5 minutes	

CIFS requirements

Item	Yes
CIFS license 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 set up and running on the 7-Mode system	
Authentication type in 7-Mode for CIFS is Active Directory (AD)	
CIFS added to the list of allowed protocols for the SVM	
DNS configured for the SVM	

Item	Yes
CIFS server configured for the SVM	
CIFS is running on the SVM	

Related concepts

[Preparing for migration](#) on page 39

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.
- 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 need to 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. Also, different projects cannot write to the same Storage Virtual Machine (SVM).

Steps

1. Log in to the tool and in the **Migrate** section, click **Get Started**.
2. From the left pane, click **New Transition Project** to create a new transition project.

Note: 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**.

Note: 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 against 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 against 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.

The hyperlink provided on the volume name, opens a dialog box that lists related qtrees and their attributes such as qtree name, owning vFiler unit, state, security, NFS export path, and NFS actual path.

Tip: It is a best practice that volumes within a single project 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**.
- 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. 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**.
10. Optional: From the **Interface Mapping** pane, provide information about the LIFs that must be created on the SVM.

If you want to...	Then...
Transition an existing 7-Mode IP address	a. Select the 7-Mode IP address. b. Click Choose Node and Port. c. From the Choose Node and Port dialog box, select the target node and port to which the LIF must be migrated and click Save. d. Click Next.

If you want to...	Then...
Create a new LIF	a. Click New LIF. b. In the Provide details for additional LIF in the destination cluster dialog box, enter the details for the new LIF and click Add. c. Click Next.

The LIFs are created on the destination SVM only if the transition completes successfully.

- 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.

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

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

- f. In the **Parameters for Transition Data Copy Operations (based on Volume SnapMirror)** pane, specify the maximum number of concurrent volume SnapMirror transfers (as a percentage of available SnapMirror transfers at run time and as a number) and the throttle limit (maximum bandwidth for all 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 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 74

[Managing SnapMirror transfers and schedule](#) on page 69

Related tasks

[Creating a data copy schedule for SnapMirror transfers](#) on page 70

[Transitioning volumes by excluding certain 7-Mode configurations](#) on page 78

[Managing logical interfaces](#) on page 79

[Removing volumes from a subproject](#) on page 80

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 validate the feasibility of transition and 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 when you have only selected 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
- LUNs
- Volumes are offline

Warning

Configurations that can cause minor problems after transition.

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:

- MetroCluster configuration
- 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 70

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. You can verify whether the copied configurations are applied properly before disconnecting the clients, thereby reducing the downtime.

About this task

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

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

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 MS 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 sub-project 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.

Related concepts

Considerations for quotas on page 91

Related tasks

Managing logical interfaces on page 79

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.

About this task

The time taken for the storage cutover to complete is only a few minutes. The time needed 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 have to be updated determines the cutover time.

The outage taken applies 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

1. From **Dashboard**, select the subproject that you want to complete.
2. 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 clients have to be disconnected from the volumes.
3. Manual: Disconnect client access.
4. Click **Complete Transition**.
Note: In case of transitioning volumes in SnapMirror relationship, secondary subproject must be completed first before completing the transition of the primary subproject.
A message is displayed warning you that the cutover operation is disruptive.
5. Click **Yes** to continue.

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

6. Click OK.

Your data is now successfully transitioned to clustered Data ONTAP.

Result

The 7-Mode IP addresses selected for the transition are unconfigured from the 7-Mode storage system and the associated LIFs created during the precutover are brought to the administrative **up** state. The 7-Mode volumes are made offline.

After you finish

Perform any post-transition steps and manually reconnect client access.

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 83

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 57

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.

About this task

- The 7-Mode Transition Tool service must be always running for the schedules to be effective.
- The data copy schedule is used for SnapMirror baseline copy, update, and resynchronization operations.
- You must create at least one data copy schedule for each 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 tool aborts the extra transfers to ensure that the tool 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, 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 Name	Days Range	Start Time	Duration	Percentage of maximum number of concurrent transfers	Maximum number of concurrent transfers	Throttle (MBps)	Update Frequency
peak_hours	Mon-Fri	09:30	8:00	50		100	0:00
dr_active	Mon-Fri	23:30	3:00		25	200	0:30
non_peak_non_dr1	Mon-Fri	17:00	6:30	75		300	1:00
non_peak_non_dr2	Mon-Fri	02:30	6:30	75		300	1:00
week_ends	Sat	02:30	53:30	90		500	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 to be used for SnapMirror baseline copy, update, and resynchronization operations, if you have other DR schedules configured or modified that requires 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, new throttle limit is considered for the next operations for this SnapMirror relationship.
- If you want the throttle limit to be effective immediately, 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, so 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

Schedule	Number of volumes and data copy state	Maximum number of concurrent SnapMirror transfers	Throttle limit	Throttle used by each transfer
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 2 volumes gets 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.

2. 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. Enter the maximum number of concurrent SnapMirror transfers to run during transition.
 - b. Enter the percentage of available streams that the tool can use for SnapMirror transfers.
By default, the tool uses 50% of the available volume SnapMirror transfers.
 - c. Enter the throttle limit to use the maximum bandwidth.
By default, the tool uses maximum throttle for configuration transition.
3. Click **Continue**.

Related tasks

[Starting baseline data copy](#) on page 63

[Creating a data copy schedule for SnapMirror transfers](#) on page 70

Transitioning volumes by excluding certain 7-Mode 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

In addition to the Interface Mapping pane, you can use Dashboard to manage the IP addresses and logical interfaces (LIFs) in a subproject. These IP addresses and LIFs are used for creating and configuring LIFs on the cluster.

Before you begin

- You must have created a project.

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

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 the volumes from the 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.

Pausing and resuming a subproject

You can pause and resume a subproject that has started.

Before you begin

- You must have started a subproject.

About this task

If you pause a subproject, the data copy operation for all the volumes in the subproject also stops.

For primary subprojects, two copy operations run simultaneously when you start the subproject. One copy operation is from the 7-Mode primary volume to the corresponding clustered Data ONTAP primary volume. The other is the from the 7-Mode primary volume to the clustered Data ONTAP secondary volume that was transitioned from the 7-Mode secondary volume. When you pause a primary subproject, only the first type of copy operation is paused. The second type of copy operation is required for the SnapMirror relationship and it continues to run as and when required. This ensures continued data protection for your SnapMirror volumes.

You can resume the subproject to continue the copy operation from where it was paused.

Steps

1. From **Dashboard**, select the subproject that you want to pause.
2. Click **Pause**.
3. To resume the paused subproject, click **Resume**.

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

Related tasks

[Starting baseline data copy](#) on page 63

Aborting a subproject

You might want to abort a running subproject for reasons such as performance of the controller or the cluster is impacted. You can abort a running subproject from Dashboard.

Before you begin

The baseline copy must have 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 also aborts the associated subproject. For example, if you abort a primary subproject, the secondary subproject is also aborted.
- If the abort 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 **Dashboard**, select the subproject that you want to abort.
2. Click **Abort**.

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.

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 50

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

- SnapMirror operations (baseline, update, or resynchronization) from the 7-Mode source volumes to the destination clustered Data ONTAP volumes are run based on the schedules you configure while creating the subproject.
The number of concurrent SnapMirror transfers that the tool uses for running the SnapMirror operations depends on the values configured with the schedules.
- You should ensure that the data copy schedules do not overlap across different subprojects, if the subprojects use the same 7-Mode controller or cluster.
- 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:
 - Create SnapMirror data copy schedules for a subproject such that it does not overlap with the existing 7-Mode SnapMirror or SnapVault schedules.

- Configure the number of concurrent SnapMirror transfers that the tool needs to use while configuring the schedules in such a way that the existing 7-Mode SnapMirror or SnapVault schedules have the required number of concurrent transfers and do not fail.
You can also release some transfers by editing the active schedule and modifying the maximum number of concurrent volume SnapMirror transfers as either zero percentage or zero.
- 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 happen.

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.

You must consider the following before using multiple paths for transition:

- If the secondary path fails before the transfer is started, the baseline transfer begins after a delay.
- After the transfer starts between the 7-Mode system and the clustered Data ONTAP system, both paths must be available for SnapMirror transfers to continue.
If any path fails when the transfer is in progress, the data copy fails.

Related concepts

[Considerations for creating a data copy schedule](#) on page 74

Related tasks

[Creating a data copy schedule for SnapMirror transfers](#) on page 70

Related information

[Documentation on the NetApp Support Site: mysupport.netapp.com](https://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 77

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 need 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:

```
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
R0 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:

```
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
```

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

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

```
[root@192.168.10.10 ]# mount 192.168.35.223:/vol/volstd10/qtrees1
transition_volume_qtreemount:192.168.35.223:/vol/volstd10/qtrees1
failed, reason
given by server: Permission denied [root@192.168.10.10 ]#
```

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

```
cluster-01::> export-policy rule create -vserver std_21 -policyname
std_2124 -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 qtree1 when the client accesses the qtree directly or through the mount point of the qtree's parent volume.

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

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 48
- 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 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>`

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 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 thin-provisioned volumes

The 7-Mode Transition Tool does not preserve thin-provision settings. If your destination aggregate does not have sufficient space to host the thick-provisioned volumes, you can use the 7-Mode Transition CLI to transition to a thin-provisioned volume in clustered Data ONTAP.

The following steps outline the high-level transition process for retaining the thin-provisioning settings:

1. From the cluster CLI, create the clustered Data ONTAP volumes equal or larger in size to the 7-Mode volumes, setting the space guarantee to **none** and type **DP**.
2. From the 7-Mode Transition Tool CLI, create a transition session and add the clustered Data ONTAP volumes created in the previous step to the session.
3. Perform data copy and complete the transition of the session.

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 can be transitioned. If the number of domains or name servers in 7-Mode exceeds the limit, DNS configuration is not 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 complete 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 complete and then run the transition start operation again.

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*.

NFS configuration settings that are transitioned

There are some common 7-Mode NFS configuration settings that are collected by the 7-Mode Transition Tool and transitioned to the clustered Data ONTAP system.

The following are some of the NFS configuration that are transitioned:

- NFS volume export rules
- NFS qtree export rules
Transitioning of qtree export rules is supported to a target cluster running Data ONTAP 8.2.1 and later.
- NFS options
- Kerberos configuration

This topic does not provide a complete list of the settings that are transitioned.

The following table shows the equivalent clustered Data ONTAP settings for some of the 7-Mode NFS options:

7-Mode configuration setting	Equivalent configuration setting in clustered Data ONTAP
nfs.mount_rootonly	<code>vserver nfs modify -vserver <i>vserver _name</i> -mount-rootonly</code>
nfs.udp.enable	<code>vserver nfs modify -vserver <i>vserver _name</i> -udp</code>
nfs.tcp.enable	<code>vserver nfs modify -vserver <i>vserver _name</i> -tcp</code>
nfs.udp.xfersize	<code>vserver nfs modify -vserver <i>vserver _name</i> -udp-max-xfer-size</code>
nfs.v3.enable	<code>vserver nfs modify -vserver <i>vserver _name</i> -v3 enable</code>
nfs.v4.enable	<code>vserver nfs modify -vserver <i>vserver _name</i> -v4.0 enable</code>
nfs.v4.read_delegation	<code>vserver nfs modify -vserver <i>vserver _name</i> -v4.0-read-delegation</code>
nfs.v4.write_delegation	<code>vserver nfs modify -vserver <i>vserver _name</i> -v4.0-write-delegation</code>
nfs.v4.id.domain	<code>vserver nfs modify -vserver <i>vserver _name</i> -v4-id-domain</code>
nfs.v4.acl.enable	<code>vserver nfs modify -vserver <i>vserver _name</i> -v4.0-acl-enabled</code>
nfs.vstorage.enable	<code>vserver nfs modify -vserver <i>vserver _name</i> -vstorage</code>
nfs.v4.snapshot.active.fs id.enable	<code>vserver nfs modify -vserver <i>vserver _name</i> -v4-fsid-change</code>

7-Mode configuration setting	Equivalent configuration setting in clustered Data ONTAP
nfs.v4.acl.max.aces	vserver nfs modify -vserver <i>vserver</i> <i>_name</i> -v4-acl-max-aces
nfs.tcp.xfersize	vserver nfs modify -vserver <i>vserver</i> <i>_name</i> -tcp-max-xfer-size
nfs.rpcsec.ctx.high	vserver nfs modify -vserver <i>vserver</i> <i>_name</i> -rpcsec-ctx-high
nfs.rpcsec.ctx.idle	vserver nfs modify -vserver <i>vserver</i> <i>_name</i> -rpcsec-ctx-idle
nfs.response.trace	vserver nfs modify -vserver <i>vserver</i> <i>_name</i> -trace-enabled
nfs.response.trigger	vserver nfs modify -vserver <i>vserver</i> <i>_name</i> -trigger

Related information

NetApp Technical Report 4067: Clustered Data ONTAP Best Practice and NFS Implementation Guide

CIFS configuration settings that are transitioned

There are some common 7-Mode CIFS configuration settings that are collected by the 7-Mode Transition Tool and transitioned to the clustered Data ONTAP system.

7-Mode Transition Tool always transitions the actual path; aliases the path that clients use to mount the source path) are ignored. Clients that use aliases must remount to the actual path after the transition.

This topic does not provide a complete list of the settings that are transitioned.

The following are some of the CIFS configuration settings that are transitioned:

- Access control lists or ACLs (either NTFS or NFS)
- CIFS preferred domain controller settings
- CIFS shares configuration and ACL configuration
- CIFS user mapping configuration
- CIFS home directory configuration
- CIFS symlinks and widelinks configuration

- Local users and groups (preserving SIDs)
Transitioning local users and groups is supported to a target cluster running Data ONTAP 8.2.1 and later.
- Audit configuration
- GPO settings
- Creating or modifying export policy rule to allow CIFS access

The following table shows the equivalent clustered Data ONTAP setting for a particular 7-Mode CIFS option:

7-Mode configuration setting	Equivalent configuration setting in clustered Data ONTAP
<code>cifs.smb2.enable</code>	<code>set -privilege advanced -confirmations off vserver cifs options modify -vserver vserver _name -smb2-enabled</code>
<code>cifs.smb2.signing.required</code>	<code>vserver cifs security modify -is-signing- required true</code>
<code>cifs.wins_servers</code>	<code>vserver cifs options modify -vserver vserver _name -wins-servers</code>
<code>cifs.grant_implicit_exe_permissions</code>	<code>vserver cifs options modify -vserver vserver _name -read-grants-exec</code>

Starting in clustered Data ONTAP 8.2.2, the `cifs.restrict_anonymous` is supported. The default value of this option in Data ONTAP operating in 7-Mode and clustered ONTAP 8.2.2 is **no_restriction**. The following table shows how this option is transitioned depending on the current values that are set in both the source and destination systems:

Value of <code>cifs.restrict_anonymous</code> option	Is transitioned?	Example
Equal on both Data ONTAP operating in 7-Mode and clustered Data ONTAP	Yes	• Clustered Data ONTAP: no_restriction • 7-Mode: no_restriction

Value of cifs.restrict_anonymous option	Is transitioned?	Example
More restrictive in Data ONTAP operating in 7-Mode	Yes, but warning issued	◦ Clustered Data ONTAP: no_restriction ◦ 7-Mode: 2
More restrictive in clustered Data ONTAP	No; warning issued	• Clustered Data ONTAP: no_restrictionnn • 7-Mode: 0

Related information

NetApp Technical Report 4191: Best Practices Guide for Clustered Data ONTAP 8.2 Windows File Services

NetApp Technical Report 3967: Deployment and Best Practices Guide for Clustered Data ONTAP 8.1 Windows File Services

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...
<code>session_name/ transition.log</code>	Debug messages that are specific to a session
<code>session_name/zapi- outbound.log</code>	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...
<code>transition-gui.log</code>	Entries of all the actions performed by using the web interface

Log file path	Contains information about...
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
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.

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.

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**.

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.

Volume goes to an inconsistent state after transition is paused

Issue

Volumes go to inconsistent state when a subproject is paused during the baseline data copy

Cause

The subproject is paused before the SnapMirror restart checkpoint is created for the volumes.

Workaround

1. Log in to the cluster.
2. Use the `snapmirror break` command to break the SnapMirror relationship.
3. Use the `snapmirror delete` command to delete the SnapMirror relationship.

4. Use the `volume delete` command to delete the clustered Data ONTAP volume.
5. Use the `volume create` command to re-create the volume.
6. Use the `snapmirror initialize` command to reinitialize the SnapMirror relationship between the 7-Mode and clustered Data ONTAP volumes.
7. Use the `snapmirror show` command with the `-snapshot-checkpoint` parameter to view the status of the baseline transfer and the restart checkpoint.

For more details about these commands, see the man pages.

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 103
- [Commands to manage volume SnapMirror relationships](#) on page 104
- [Commands to manage credentials of Data ONTAP systems](#) on page 105
- [Commands to manage the volume pairs in a session](#) on page 105
- [Commands to manage LIF details](#) on page 105
- [Commands to manage session properties](#) on page 105
- [Commands to manage transition jobs](#) on page 106
- [Commands to manage transition schedules](#) on page 106

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>

If you want to...	Use this command...
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>
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>

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>
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>

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 50

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.

Copyright information

Copyright © 1994–2014 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, 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, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. 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.

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

7-Mode

- data migration to clustered Data ONTAP [6](#)
- preparing for assessment [19](#)
- preparing for transition [44](#)
- transitioning volumes and configuration to clustered Data ONTAP [6](#)

7-Mode configurations

- applying [65](#)
- copying [65](#)
- testing [65](#)
- verifying [65](#)

7-Mode Transition Tool

- configuration file [13](#)
- installing [11](#), [12](#)
- log files [98](#)
- logging in using the web interface [14](#)
- network port requirements [46](#)
- system requirements [11](#)
- uninstalling [11](#), [16](#)
- upgrading [11](#)
- web interface
 - logging in to 7-Mode Transition Tool [14](#)

A

- adding systems for migration [56](#)
- applications
 - supported versions for transition assessment [17](#)
- assessment
 - host and storage inventory information [17](#)
 - how to use the executive summary for transition [24](#)
 - preparing for transition [17](#)
 - preparing the 7-Mode system [19](#)
 - when to perform while transitioning [6](#)
- assessment limits
 - maximum number of controllers, hosts [9](#)
- assessment reports
 - adding systems to create [21](#)

B

- backup of volumes in transition
 - using dump [93](#)
- baseline data copy
 - data copy from 7-Mode

starting [63](#)

starting [63](#)

transition phase described [26](#)

best practices

during data copy [74](#)

for data copy [74](#)

browsers

supported versions for the web interface [11](#)

C

CIFS

prerequisites to transition [48](#)

transitioned settings [95](#)

CIFS migration

workflow [37](#)

client access

volumes in data protection relationship [31](#)

clustered Data ONTAP

transitioning data and configuration from 7-Mode [6](#)

clusters

adding [56](#)

preparing for transition [41](#)

collection

host and storage inventory information [17](#)

commands to manage

credentials [103](#)

jobs [103](#)

LIFs [103](#)

session properties [103](#)

SnapMirror relationships [103](#)

transition operations [103](#)

transition sessions [103](#)

volume pairs [103](#)

comments

how to send feedback about documentation [111](#)

compression

considerations for transition [90](#)

configuration

not supported for transition [7](#)

configuration file

modifying [13](#)

configuration transition phase

transition of stand-alone volumes [28](#)

considerations

for quotas while transitioning [91](#)

- for transition thin-provisioned volumes [91](#)
 - for transitioning DNS and LDAP [92](#)
 - for transitioning FlexClone volumes [91](#)
- considerations for NFS [85](#)
- considerations for NFS exports transition [85](#)
- considerations for NFS transition [85](#)
- controllers
 - adding [56](#)
 - assessing transition readiness of [20](#)
- credentials
 - commands to manage [103](#)
- cutover
 - guidelines for performing [85](#)
- cutover phase
 - volumes in a data protection relationship [31](#)

D

- dashboard
 - aborting subprojects [81](#)
 - prechecks [61](#)
- data copy
 - considerations for scheduling [74](#)
 - priority for transfers [74](#)
- data copy phase
 - transition of stand-alone volumes [28](#)
 - volumes in a data protection relationship [31](#)
- data copy schedule
 - creating [70](#)
 - deleting [75](#)
 - editing [75](#)
 - managing [69](#)
- data copy update
 - scheduling [70](#)
- Data ONTAP
 - 7-Mode versions supported [39](#)
 - 7-Mode versions supported for transition assessment [17](#)
- data resynchronization operation
 - scheduling [70](#)
- deduplication
 - considerations for transition [90](#)
- DNS
 - considerations for transition [92](#)
- documentation
 - how to send feedback about [111](#)
- downloading
 - the log files [98](#)
- dump support for transition
 - understanding [93](#)

E

- editing subprojects
 - using web interface [69](#)
- error messages
 - 7-Mode Transition Tool installer is blocked [100](#)
- executive summary
 - how you can use it for transition assessment [24](#)

F

- features
 - not supported for transition [7](#)
- feedback
 - how to send comments about documentation [111](#)
- files
 - collecting the logs [98](#)
- FlexClone volumes
 - considerations for transition [91](#)
- flowcharts
 - CIFS migration [37](#)
 - NFS migration [36](#)

G

- giveback
 - impact on transition [92](#)
- guidelines
 - for creating data copy schedules [74](#)

H

- hosts
 - assessing transition readiness of [20](#)
- hosts, supported
 - versions for transition assessment [17](#)

I

- importing
 - inventory report XML file for transition assessment [23](#)
- information
 - how to send feedback about improving documentation [111](#)
- installation
 - workaround for error message during [100](#)
- installing
 - 7-Mode Transition Tool [11](#), [12](#)

- inventory information
 - collecting and assessing [17](#)
- inventory reports
 - importing for transition assessment [23](#)
- issues during pause operation [101](#)

J

- jobs
 - commands to manage [103](#)

L

- LDAP
 - considerations for transition [92](#)
- licenses
 - requirements for transition [39](#)
- LIF
 - See* logical interfaces
- LIFs
 - commands to manage [103](#)
- limits
 - number of controllers and hosts for assessment [9](#)
- log files
 - 7-Mode Transition Tool [98](#)
- logging in
 - to the 7-Mode Transition Tool [14](#)
- logical interfaces
 - managing [79](#)
 - See also* LIFs
- logs
 - collecting 7-Mode transition Tool logs [98](#)

M

- migrate
 - Data ONTAP version requirements [39](#)
 - supported 7-Mode versions [39](#)
- migration
 - downloading the log files [98](#)
 - preparing for [39](#)
 - when to perform while transitioning [6](#)
 - workflow for CIFS [37](#)
 - workflow for NFS [36](#)
- migration phases
 - baseline data copy [26](#)
 - precutover [26](#)
 - preparation [26](#)
 - storage cutover [26](#)
- multiple paths

- considerations for data copy using [83](#)

N

- network ports
 - requirements [46](#)
- NFS
 - prerequisites to transition [47](#)
 - transitioned settings [93](#)
- NFS migration
 - workflow [36](#)
- NFS only transition
 - NFS only configuration transition [85](#)

O

- on-demand SnapMirror updates
 - performing [77](#)
- operational changes
 - when to perform while transitioning [6](#)
- overview
 - transition [6](#)

P

- phases
 - stand-alone volumes, transition [28](#)
 - volumes in data protection relationship, transition [31](#)
- planning
 - when to perform while transitioning [6](#)
- ports
 - requirements [46](#)
- post transition steps
 - transition of stand-alone volumes [28](#)
- prechecks
 - error [62](#)
 - information [62](#)
 - notification [62](#)
 - running [61](#)
 - severity levels [62](#)
 - warning [62](#)
- precutover
 - transition phase described [26](#)
- precutover phase
 - volumes in a data protection relationship [31](#)
- preparation
 - for transition, checklist [51](#)
 - transition phase described [26](#)
- preparation phase
 - transition of stand-alone volumes [28](#)

- volumes in a data protection relationship [31](#)
- preparing
 - 7-Mode system for assessment [19](#)
 - 7-Mode system for transition [44](#)
 - cluster for transition [41](#)
 - for migration [39](#)
- protocols
 - preparing the 7-Mode system for assessment [19](#)

Q

- quotas
 - considerations for transition [91](#)

R

- removing
 - volume from subproject [80](#)
- reports
 - creating assessment report to check transition readiness [21](#)
 - importing inventory reports for transition assessment [23](#)
- requirements
 - 7-Mode system, hosts, and host application versions supported for transition assessment [17](#)
 - 7-Mode Transition Tool [46](#)
 - for the tool, system [11](#)
 - for transitioning CIFS configurations [48](#)
 - for transitioning NFS configurations [47](#)
 - licenses for transition [39](#)
- requirements for transition
 - transition service [40](#)
- restrictions
 - on transition [40](#)

S

- SAN
 - unsupported features [7](#)
- schedules
 - guidelines for creating data copy [74](#)
- secondary subproject
 - aborting [81](#)
- settings
 - CIFS, supported for transition [95](#)
 - NFS, supported for transition [93](#)
- SnapMirror baseline copy
 - scheduling [70](#)
- SnapMirror relationships

- commands to manage [103](#)
 - considerations for transition [83](#)
- SnapMirror schedule
 - creating [70](#)
- SnapMirror schedules
 - managing [69](#)
- SnapMirror transfer schedule
 - deleting [75](#)
 - editing [75](#)
 - modifying [75](#)
- SnapMirror updates
 - performing on-demand [77](#)
- stand-alone subproject
 - aborting [81](#)
- storage cutover
 - transition phase described [26](#)
- storage efficiency
 - considerations for transition [90](#)
- subprojects
 - aborting [81](#)
 - adding volumes [69](#)
 - editing [69](#)
 - managing [69](#)
 - modifying transition configuration [69](#)
 - pausing and resuming [81](#)
 - removing volumes [69](#)
- suggestions
 - how to send feedback about documentation [111](#)
- system requirements
 - for the tool [11](#)

T

- takeover
 - impact on transition [92](#)
- thin provisioning
 - retaining settings for transition [91](#)
- throttle limit
 - specifying [70](#)
- throttle limits
 - specifying [77](#)
- transfers
 - priority for [74](#)
 - transition considerations for concurrent SnapMirror [83](#)
- transition
 - assessing controllers and hosts are ready for [20](#)
 - assessing hosts and controllers for [17](#)
 - best practices during the different phases of [26](#)
 - checklist [51](#)

- CIFS configuration prerequisites [48](#)
 - collecting inventory and creating report for assessment [21](#)
 - commands to manage [103](#)
 - configuration file for [13](#)
 - considerations for deduplication and compression [90](#)
 - considerations for FlexClone volumes [91](#)
 - considerations for local users and groups [89](#)
 - considerations for quotas [91](#)
 - considerations for SnapMirror [83](#)
 - considerations for thin provisioning [91](#)
 - downloading the log files [98](#)
 - guidelines for cutover [85](#)
 - how to use the execute summary for assessment [24](#)
 - impact of takeover and giveback [92](#)
 - importing inventory report XML file for readiness check [23](#)
 - license requirements [39](#)
 - local users and groups
 - considerations during CIFS transition [89](#)
 - NFS configuration prerequisites [47](#)
 - overview [6](#)
 - phases of [31](#)
 - prechecks [62](#)
 - preparing the 7-Mode system [44](#)
 - preparing the cluster for [41](#)
 - project definition [107](#)
 - restrictions on [40](#)
 - session definition [107](#)
 - stand-alone volumes [28](#)
 - subproject definition [107](#)
 - terminology [107](#)
 - unsupported features [7](#)
 - volumes in a data protection relationship [31](#)
 - transition assessment
 - 7-Mode system, hosts, and host application versions supported for [17](#)
 - transition complete phase
 - transition of stand-alone volumes [28](#)
 - transition considerations
 - DNS and LDAP [92](#)
 - Snapshot copies [91](#)
 - transition issues
 - cannot select a secondary volume [99](#)
 - transition limits
 - number of active projects [9](#)
 - number of volumes in a subproject [9](#)
 - transition operations
 - commands to manage [103](#)
 - transition project
 - completing [67](#)
 - creating [57](#)
 - deleting [82](#)
 - managing [69](#)
 - running [50](#)
 - workflow [50](#)
 - transition projects
 - prechecks [61](#)
 - transition sessions
 - commands to manage [103](#)
 - transition subprojects
 - completing [67](#)
 - pausing [81](#)
 - resuming [81](#)
 - transition-tool.conf
 - modifying [13](#)
 - transitioned settings
 - CIFS [95](#)
 - NFS [93](#)
 - transitioning
 - from 7-Mode to clustered Data ONTAP [6](#)
 - transitioning volumes
 - excluding CIFS local users and groups [78](#)
 - excluding configurations [78](#)
 - netgroup configurations [78](#)
 - UNIX user and groups [78](#)
- ## U
- uninstalling
 - 7-Mode Transition Tool [16](#)
 - unsupported features
 - for transition [7](#)
 - updates
 - performing on-demand SnapMirror [77](#)
 - upgrading
 - 7-Mode Transition Tool [15](#)
- ## V
- volume goes to an inconsistent state
 - during pause [101](#)
 - volume mapping
 - no aggregates assigned [100](#)
 - volume pairs
 - commands to manage [103](#)
 - volumes
 - adding [57](#)
 - LIFs
 - configuring [57](#)

- logical interface
 - configuring [57](#)
- volumes in transition
 - backing up [93](#)
- Vservers
 - See* SVMs

W

- web interface

- troubleshooting [98](#)
- workflows
 - CIFS migration [37](#)
 - NFS migration [36](#)