



Copy-free transition

ONTAP 7-Mode Transition

NetApp
February 11, 2024

This PDF was generated from https://docs.netapp.com/us-en/ontap-7mode-transition/copy-free/concept_copy_free_transition_using_7mtt.html on February 11, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Copy-Free Transition Guide 1
 - Transition overview 1
 - Collecting and assessing the inventory information 4
 - Copy-free transition workflow 14
 - Preparing for copy-free transition 18
 - Transitioning 7-Mode aggregates using copy-free transition 57
 - Completing the transition 79
 - Transitioning a SnapMirror relationship 86
 - Troubleshooting transition issues 89
 - Performing a transition rollback to 7-Mode 96

Copy-Free Transition Guide

This guide describes how to transition from a 7-Mode HA pair to an ONTAP cluster using the 7-Mode Transition Tool, without having to copy data from disks; the existing disk shelves are connected to the new cluster.

Transition overview

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

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

[7-Mode Transition Tool Release Notes](#)

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

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

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

Related information

[NetApp Documentation: ONTAP 9](#)

Copy-free transition by using 7-Mode Transition Tool

The 7-Mode Transition Tool enables you to collect inventory and to assess 7-Mode controllers, hosts, switches, and applications for transition. After assessment, you can use a copy-free transition to migrate your data and configurations to ONTAP. In copy-free transition, you disconnect all the disk shelves from a 7-Mode HA pair and connect them to an HA pair in the target cluster.

Copy-free transition significantly reduces the migration cost by enabling the reuse of 7-Mode disk shelves. The overall duration for performing the transition is faster because data copy is not required.


The unit of a copy-free transition is an HA pair. You must move all the disk shelves from the 7-Mode HA pair to

the target cluster nodes.

The metadata of the 7-Mode aggregates and volumes is converted to the ONTAP format by the 7-Mode Transition Tool. The time taken for this conversion does not depend on the size of the aggregates and volumes. For example, the time taken to convert a 10 GB aggregate to the ONTAP format is the same as the time required to convert a 100 TB aggregate.

Copy-free transition involves a disruption to data access. However, the total time taken to perform the data migration is faster because no data copy is required.

The following illustration shows the before and after scenarios for copy-free transition from a 7-Mode HA pair to a two-node cluster:



The 7-Mode Transition Tool runs on a Windows system and provides web interface for managing transition operations.

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

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

- Collecting inventory information from ONTAP systems (7-Mode controllers and nodes in the cluster), hosts, switches, and host applications.
- Generating the FC Zone plan to configure the zones for SAN FC transition.
- Assessing the features and functionalities of the 7-Mode systems, and identify how these features and functionalities work in the ONTAP version selected for the transition.

Moving data and configurations from 7-Mode to ONTAP

Copy-free transition reuses the 7-Mode disk shelves to copy 7-Mode data from 7-Mode to ONTAP. You can perform the following tasks by using the 7-Mode Transition Tool for copy-free migration:

- Planning your transition to map the 7-Mode controllers or vFiler units to the target SVMs, and design the namespace.
- Running prechecks to verify the compatibility of the 7-Mode systems and target cluster nodes for transition.
- Importing 7-Mode disk shelves in the following ways:
 - From a 7-Mode HA pair to a new HA pair in a new cluster
 - From a 7-Mode HA pair to a new HA pair in an existing cluster that has additional data-serving nodes
 - From a 7-Mode HA pair to an HA pair that has data aggregates in an existing cluster that is serving data
 - From an HA pair that contains volumes in a volume SnapMirror relationship to an HA pair in a new or existing cluster

You must manually create the cluster peer relationship after transition; however, a rebaseline transfer is not required, and you can retain the SnapMirror relationship after transition.

- Transitioning 7-Mode configurations to SVMs.

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

- Rolling back storage and configurations to 7-Mode if transition to ONTAP fails.

The tool generates the list of steps that are required to roll back to 7-Mode. You must manually perform these rollback steps on the 7-Mode systems and the cluster.

Transition terminology

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

- **Collect and assess**

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

- **Migrate**

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

- **Project**

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

- **Project group**

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

Collecting and assessing the inventory information


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

You can collect inventory information in two ways:

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

In both cases, you must use the latest version of the Inventory Collect Tool to collect the inventory.

To assess the inventory information for a copy-free transition, you must select both of the nodes of the source 7-Mode HA pair. Although the assessment is done on a per-node basis, if a single node is not qualified for transition, the entire HA pair cannot be transitioned.



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

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

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

[NetApp Interoperability Matrix Tool](#)

Preparing the 7-Mode systems and hosts for transition assessment

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

Steps

1. Enable HTTPS on the 7-Mode system:

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

2. Enable TLS on the 7-Mode system:

```
options tls.enable on
```



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

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

- a. Set up and start SSL:

```
secureadmin setup ssl
```

- b. Enable SSL:

```
options ssl.enable on
```

- c. Disable SSLv2 and SSLv3:

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

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



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

4. Enable SSH on the 7-Mode system:

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

```
secureadmin setup -f ssh
```

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

- b. Enable SSH:

```
secureadmin enable ssh2
```

- c. Enable password authentication on the SSH server:

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

- d. Enable SSH access to the host:

```
options ssh.access
```

5. Prepare your Windows host systems:

- Enable WMI access.

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

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

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

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

```
SystemName\Username
```

6. Enable SSH on the Linux or ESXi host.

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

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

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

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

Related information

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

Assessing controllers and hosts

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

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

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

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

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

You can collect inventory information for controllers, hosts, and FC switches by adding the systems to the 7-Mode Transition Tool. You can then create an assessment report to

assess the features and functionalities of those systems, and to identify how they work in the ONTAP version selected for transition.

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

The readme file is located at `_7-Mode_Transition_Tool_installed_location\bin\ict`.

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

Each system details must be in the following format:

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

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

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

Steps

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

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

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

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

iii. Select the system type:

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

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

5. Click **Add**.

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

6. Generate the transition assessment report:

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

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

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

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

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

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

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

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

Generating an assessment report by importing the inventory report XML

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

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



To assess hosts and controllers for copy-free transition, you must use the latest version of the Inventory Collect Tool to collect inventory.

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

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

Steps

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

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

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

The `AssessmentWorkbook` and `AssessmentExecutiveSummary` reports are generated in XML format.

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

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

Generating an FC zone plan

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

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

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

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

Steps

1. In the Collect and Assess section, click **Add Systems**.

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

3. Click **Add**.

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

4. Generate the transition assessment report with the FC zone plan:
 - a. Select the systems, including the required FC switches, for transition assessment.
 - b. Click **Create Transition Assessment Report**.
 - c. In the pair systems and start FC zone planner dialog box, select the 7-Mode systems (single controller or HA pair), the switch connected to the cluster, and an SVM in the cluster.

If you plan to consolidate the FC LUNs to a single SVM by rehosting transitioned volumes, select the 7-Mode HA pair and the target SVM.



It is a best practice to consolidate the FC LUNs to a single SVM for preserving the 7-Mode single-system image (SSI) configuration.

If you do not plan to consolidate the FC LUNs, you must generate the FC zone plan for each 7-Mode controller and the corresponding target SVM.

- d. Click **FC Zoning for paired systems**.
- e. In the Create Transition Assessment Report dialog box, select the Data ONTAP version of the target cluster.
- f. Specify a prefix for the file name of the reports.
- g. Click **Generate Report**.

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

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


Supported configurations for generating an FC zone plan

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

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

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

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



How you can use the assessment executive summary for transition assessment

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

The executive summary has the following main sections:

Target cluster

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

Data collection summary

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

Transition feasibility and recommended transition methodology

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

Copy-free transition feasibility: This section lists the number of controller-level prechecks that resulted in errors and warnings for copy-free transition. If a precheck fails for any one controller in the HA pair, you cannot transition the HA pair by using copy-free transition. You must resolve all errors and warnings before transitioning the HA pair. Details of these prechecks are available in the CFT Precheck Summary tab of the assessment workbook.

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

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

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

Storage inventory

This section provides the following information:

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

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

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


Data collection errors

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

again.

Copy-free transition workflow

The copy-free transition workflow includes preparing for the transition, performing the transition, and completing the transition. Some of these tasks must be done manually on the 7-Mode systems and the cluster.



Manual: Tasks that cannot be done from 7-Mode Transition Tool

Phases of copy-free transition

Copy-free transition using the 7-Mode Transition Tool consists of the following phases: planning, SVM provisioning, exporting and halting, cabling, importing, preproduction testing, starting production, and committing. You should understand the phases to manage the transition effectively.


Copy-free transition is a disruptive operation. Therefore, you must plan for the downtime for the applications and workload running on the 7-Mode storage systems.

In most cases, the storage cutover time can be 3 through 8 hours. Cutover time includes the time taken by the tool to perform two automated operations—the export and halt operation and the import operation—as well as the time taken for manually cabling the disk shelves to the new controllers. The export and halt operation and the import operation together can up to 2 hours.



For scaled configurations, the export and halt operation and the import operation together can take more than 2 hours. 7-Mode Transition Tool detects such conditions and provides a warning.

Cabling of the disk shelves can take from 1 hour through 6 hours. This cutover time guidance does not include the time for the required preproduction testing and assumes an error-free transition without unexpected failures such as disk failure.



Planning the project

You can plan the following details about the source and target of a copy-free transition project:

- 7-Mode HA pair and vFile unit details
- Target cluster nodes and mapping of source controllers to target nodes
- 7-Mode controller or vFile unit to SVM mapping
- IP addresses to transition (new LIFs or existing 7-Mode IP addresses) and the IPspaces and broadcast domains on the SVM



The 7-Mode Transition Tool does not support the transition of FC and iSCSI LIFs. These LIFs must be manually configured on the SVMs before transition.

In this phase, prechecks are run to verify whether the 7-Mode HA pair is ready to be migrated to clustered Data ONTAP. The 7-Mode Transition Tool also verifies that the cluster is configured properly and can support the transition.

You must resolve any errors 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 transition. You can run the prechecks multiple times to verify that all the errors have been resolved.

SVM provisioning

After planning your transition project, you must perform some manual tasks, such as adding licenses, creating the CIFS server, and creating SAN LIFs, to prepare the cluster and SVMs for transition.

You can then apply the configurations on the SVMs by using the tool. All the 7-Mode controller or vFiler unit level configurations are transitioned to the mapped SVM. Volume and LUN configurations are not transitioned during this phase; they are transitioned in the import phase.

At the end of this phase, you should manually verify the configurations applied to SVMs and make the necessary changes.

Exporting storage configurations and halting 7-Mode systems

This phase starts the cutover window for copy-free transition. Client access must be manually disconnected. However, all the NAS and SAN services must be up and running on the 7-Mode HA pair. This is because the 7-Mode Transition Tool requires all services to be up and running for collecting the volume-level configurations from the 7-Mode systems.

The tool performs the following operations in the export phase:

- Collects all volume and storage configurations
- Creates a Snapshot copy of each transitioning aggregate

This Snapshot copy is used for rolling back to 7-Mode, if required.

- Boots the 7-Mode controllers in maintenance mode
- Removes disk ownership from the disks attached to the 7-Mode controllers
- Disables disk autoassignment on the target cluster nodes

Cabling the 7-Mode disk shelves

You must perform the tasks in this phase manually. You must ensure that the disk shelf IDs are unique across the 7-Mode controllers and target cluster nodes.



If there are duplicate shelf IDs, you must change the disk shelf IDs and power cycle the disk shelves.

You must disconnect all the 7-Mode disk shelves and hot-add them to the target cluster nodes. After the disk shelves are connected to the target cluster nodes, you must power cycle the disk shelves.

It is a best practice to manually verify the cabling by using Config Advisor. Config Advisor is a configuration validation and health check tool for NetApp systems. It can be deployed at both secure sites and non-secure sites for data collection and system analysis.

You can then verify the cabling by using the 7-Mode Transition Tool to proceed with the transition. The 7-Mode Transition Tool performs only a subset of the cabling checks that are performed by Config Advisor.

Importing 7-Mode data and configurations

All the storage objects (aggregates, volumes, and LUNs) and the associated configurations are transitioned during this phase.

The tool performs the following operations in the import phase:

- 7-Mode disks are assigned to the mapped target cluster nodes.
- All 7-Mode aggregates, volumes, and LUNs are converted to the clustered Data ONTAP format.
- LIFs are configured on the SVMs in the administrative up state.
- All volume-level and LUN-level configurations are applied.

Preproduction testing

You must manually test all the transitioned aggregates, volumes, and configurations that are applied to the target SVMs during this phase. You must also perform all manual tasks for completing your configuration—for example, configuring hosts and performing host remediation for SAN hosts.

You cannot perform certain operations on the transitioned aggregates or volumes during this phase. There are also certain operations that are not recommended during the testing phase. This is to ensure a successful rollback operation in case you decide to roll back to 7-Mode.

You must also manually test all applications and workloads thoroughly before starting data access in a production environment.



The aggregates might run out of space because of the aggregate Snapshot copies and the write operations that are performed during the testing. If the free physical space is less than 5% of the total space, the aggregates are taken offline. You must regularly monitor the free physical space available in the transitioned aggregates to avoid space issues.

Starting production

After testing all workloads and applications, you can start client access to the transitioned data in the production environment. This stage of transition—where production is started but the project is not yet committed—is the final stage of transition when you can decide to roll back to 7-Mode. You must not prolong this phase because of the following reasons:

- The probability of running out of space in the transitioned aggregates increases as new data is written to the volumes.
- Any new data written to the volumes during this stage will not be available after rollback.

Committing the project

In this final stage of transition, the aggregate-level Snapshot copies that were created during the export phase are deleted.


You cannot roll back to 7-Mode after you commit the 7-Mode aggregates and complete the transition.

Related information

[NetApp Downloads: Config Advisor](#)

Preparing for copy-free transition

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



Related information

Requirements for copy-free transition

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

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

7-Mode Transition Tool Release Notes

- **Platform models**

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

- **Data ONTAP in 7-Mode source systems**

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

- **ONTAP target systems**

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

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

- **HA configuration**

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

- **Disk shelf models**

The following disk shelf models are supported:

- DS4486
- DS4246
- DS4243



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

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



The disk shelf model DS14mk2 FC is not supported.

- **Disk firmware**

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

[NetApp Downloads: Disk Qualification Package](#)

[NetApp Downloads: Disk Drive Firmware](#)

[NetApp Downloads: Disk Shelf Firmware](#)

- **Tool to verify cabling**

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

[NetApp Downloads: Config Advisor](#)

Tools and documentation required for copy-free transition

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

Config Advisor

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

[NetApp Downloads: Config Advisor](#)

Documentation

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

[SAN host transition and remediation](#)

Port requirements for communicating with the 7-Mode Transition Tool

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

Ports that must be open on the 7-Mode systems

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

Ports that must be open on the cluster

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

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

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

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

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

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



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

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



You must verify that firewalls do not block this traffic.

Related information

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

Preparing the 7-Mode HA pair for transition

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

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

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

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

HTTPS is enabled by default.

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

```
secureadmin setup ssl
```

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

- b. Enable SSL:

```
options ssl.enable on
```

This option must be enabled for allowing communication over TLS.

- c. Enable TLS:

```
options tls.enable on
```

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

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

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

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

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

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

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

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

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

Steps

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

```
sp setup
```



```

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

```

- b. Verify the SP network configuration settings:

sp status

```

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

```

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

options sp.ssh.access host=7mtt_host

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



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

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

ssh username@SP_IP_address

When prompted, enter the password for the user name.

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

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

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

rlm setup

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

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

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

rlm status

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

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

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

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



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

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

```
ssh username@RLM_IP_address
```

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

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

Preparing the network for transition

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

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

Steps

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

```
network port vlan create
```

or

```
network port ifgrp create
```

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

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

```
network ipspace create
```

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



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

Related information

[Network and LIF management](#)

Considerations for transitioning 7-Mode IP addresses

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

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

Preparing the cluster for transition

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

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

Software setup

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

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

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



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

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

```
ssh username@cluster_mgmt_IP
```

2. Enable SSLv3 or FIPS on the cluster:

If you want to enable...	Enter...
SSLv3	system services web modify -sslvs3 -enabled true
FIPS 140-2 compliance	system services web modify -ssl-fips -enabled true

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



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

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

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

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

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

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

system services firewall policy show -policy mgmt

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

System administration

Gathering cabling information for transition

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

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

Copy-free transition cabling worksheet

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

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

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

sysconfig slot_number

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

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

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

system node run -node node_name -command sysconfig -a

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

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

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

[NetApp Hardware Universe](#)

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

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

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

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

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

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

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

Related information

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

[DiskShelf14mk2 AT Hardware Service Guide](#)

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

Copy-free transition cabling worksheet

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

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

Preparing 7-Mode aggregates and volumes for transition

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

Restrictions for transitioning 7-Mode aggregates and volumes

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

Volume types

The following types of volumes are not supported for transition:

- Traditional volumes

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

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

- SnapLock volumes

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

- FlexCache volumes

Aggregate and volume states

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

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

FlexClone volumes

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

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

Volume with qtrees that belong to a different vFiler unit

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

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

Inode to parent pathname translation setting

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

```
vol options vol_name no_i2p off
```

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

Preparing for transitioning 7-Mode systems with 32-bit aggregates

32-bit aggregates, volumes, and Snapshot copies are not supported in ONTAP 8.3 and later. Therefore, you must expand the 32-bit aggregates to 64-bit, and then find and remove any 32-bit volumes and Snapshot copies from the 7-Mode system before transition.

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

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

Related information

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

Expanding an aggregate to the 64-bit format

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

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

For volumes in a SnapMirror relationship, the destination volume inherits the format of the source volume while the mirror is intact. If the aggregate you are expanding contains a destination volume whose source is a 32-bit volume and you break the mirror before expanding the aggregate, the destination volume is expanded to the 64-bit format. However, if you reestablish the mirror and the source volume is still 32-bit, the destination volume returns to the 32-bit format. For this reason, you must expand the aggregate containing the source volume before reestablishing the SnapMirror relationship if you want to expand all 32-bit volumes in the aggregate to the 64-bit format.

Steps

1. Enter advanced privilege mode:

```
priv set advanced
```

2. Initiate the expansion:

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

3. Perform the appropriate action:

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

4. Display the status of the expansion:

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

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

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

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

6. Return to administrative privilege mode: **priv set admin**

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

Finding and removing 32-bit volumes and Snapshot copies

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

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

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

Steps

1. Enter advanced mode:

```
priv set advanced
```

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

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

Each volume in the aggregate is displayed with its format.

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

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

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

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

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

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

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



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

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

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

6. Return to the administrative privilege level:

```
priv set admin
```

Aggregate space requirements for transition

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

Physical space in the aggregates

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

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

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

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

Logical space in the aggregates

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

Snapshot spill

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

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

Space guarantee settings

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

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

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

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

```
vol options volume_name guarantee volume
```

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

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

- **File-guaranteed volumes**

File guarantee is not supported in ONTAP.

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

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

```
vol options volume_name guarantee volume
```

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

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

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

During the transition, if these volumes contain space-reserved LUNs or files, their space guarantee will be automatically converted to `volume`, but the space guarantee will be disabled initially. You must create sufficient free space on the aggregates and then manually enable the guarantee by using the following

command:

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

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

Additional consideration for Flash Pool aggregates

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

Related information

[Ignorable errors during transition](#)

[Disk and aggregate management](#)

Preparing to transition name services

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

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

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

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

Configurations that are transitioned

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

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

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

Unsupported configurations in ONTAP

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

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

Configurations that must be manually configured

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

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

Related information

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

[NFS management](#)

[Network and LIF management](#)

Considerations for transitioning DNS, NIS, and LDAP configurations

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

Considerations for DNS transition

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



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

Considerations for NIS transition

- The length of the NIS domain name on the 7-Mode system must not exceed 64 characters.

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

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

Considerations for LDAP transition

- If separate base values and scope values are specified for user mapping (`ldap.usermap.base`) and user password (`ldap.base.passwd`) lookups in the 7-Mode system, the base values and scope values for only the user password are transitioned.

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

Considerations for transitioning netgroups and UNIX users and groups

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

Considerations for netgroups

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

- Exclude the transition of netgroups.

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

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

Considerations for UNIX users and groups

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

- Exclude the transition of UNIX users and groups.

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

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

Related information

[NFS management](#)

Preparing for NFS transition

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

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

Prerequisites for transitioning NFS configurations

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

7-Mode prerequisites

- NFS must be licensed.
- If MultiStore is licensed, NFS must be enabled on all of the vFiler units.
- NFS service must be running on the 7-Mode systems during transition.

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

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

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

Cluster prerequisites

- NFS must be licensed.

Related information

[NetApp Documentation: ONTAP 9](#)

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

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

Supported configurations for transition

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

- NFS options:

- `nfs.udp.xfersize`
- `nfs.v4.id.domain`
- `nfs.v4.acl.max.aces`
- `nfs.tcp.xfersize`
- `nfs.rpcsec.ctx.high`
- `nfs.rpcsec.ctx.idle`
- `nfs.response.trigger`
- `wafl.default_nt_user`
- `nfs.mount_rootonly`
- `nfs.tcp.enable`
- `nfs.udp.enable`
- `nfs.response.trace`
- `nfs.v4.read_delegation`
- `nfs.v4.write_delegation`
- `nfs.v4.acl.enable`
- `nfs.vstorage.enable`
- `nfs.v3.enable`
- `nfs.v4.enable`

- **NFS export rule:**

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

- **Export rules with Kerberos security krb5p**

See the precheck results for details about these NFS configurations.

Unsupported configurations in ONTAP

The following NFS configurations are not supported in ONTAP:

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

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

Configurations that must be manually transitioned

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

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

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

Configurations that are functionally different in ONTAP

The following NFS configurations are functionally different in ONTAP:

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

Related information

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

[NFS management](#)

How NFS exports are transitioned

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

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

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

volume of the SVM with the new export policy.

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

- When 7-Mode volumes with export configurations that are not supported in ONTAP are transitioned, these volumes are exported to disallow access to all NFS clients.

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

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

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

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

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

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

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

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

Consider the following export rule configured in the 7-Mode storage system (192.168.26.18) that allows read/write access to the volume volstd10 and qtree qtree1 for the NFS client 192.168.10.10:

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

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

```

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

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

cluster-01::>

```

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

```

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

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

cluster-01::>

```

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

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

```
[root@192.168.10.10 ]# mount 192.168.35.223:/vol/volstd10/mtree1
transition_volume_qtreemount:192.168.35.223:/vol/volstd10/mtree1 failed,
reason
given by server: Permission denied [root@192.168.10.10 ]#
```

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

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

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

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

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

cluster-01::>
```

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

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

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

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

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

```
[root@192.168.10.10]# mount 192.168.26.18:/vol/volstd10 transition_volume  
[root@192.168.10.10]# cd transition_volume/qtree1  
[root@192.168.10.10]# ls transition_volume/qtree1  
[root@192.168.10.10]# mkdir new_folder  
[root@192.168.10.10]# ls  
new_folder  
[root@192.168.10.10]#
```

In ONTAP, the NFS client 192.168.10.10 has only read-only access to the qtree qtree1 when the client accesses the qtree directly or through the mount point of the qtree's parent volume.

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

Related information

[NFS management](#)

Preparing for SMB/CIFS transition

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

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

Prerequisites for transitioning CIFS configurations

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

met, the tool does not transition the configuration.

7-Mode prerequisites

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

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

- The authentication type for CIFS must be Active Directory (AD) or Workgroup.

Cluster prerequisites

- The CIFS license must be added.
- The following CIFS authentication methods are supported in different ONTAP versions:
 - Clustered Data ONTAP 8.2.x and 8.3.x support AD authentication.
 - ONTAP 9.0 or later supports AD authentication and Workgroup authentication.
- The following table identifies which authentication method must be used on the target SVM:

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

- You can transition the CIFS configuration from 7-Mode to ONTAP if the AD domains do not match between the 7-Mode CIFS server and the target SVM CIFS server.

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

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

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

If you want to...	Do the following...
Transfer or preserve the CIFS server identity to the target SVM	<div data-bbox="844 159 1481 338" style="border: 1px solid #ccc; border-radius: 5px; padding: 10px; margin-bottom: 10px;"> <p>You have the following two options to create the CIFS server:</p> </div> <p>a. Applicable for all versions of ONTAP:</p> <ul style="list-style-type: none"> ◦ Before the SVM provision phase, you must reconfigure the CIFS server on the 7-Mode system by using a temporary CIFS identity. <p>This reconfiguration allows the original CIFS server identity to be configured on the SVM. You must verify that the CIFS server is running on the 7-Mode system during the “SVM Provision” and “Export & Halt” phases with the new temporary identity. This action is required to read CIFS configurations from 7-Mode during the SVM Provision and “Export & Halt” phases.</p> <ul style="list-style-type: none"> ◦ You must configure the CIFS server on the target SVM with the original 7-Mode CIFS identity. ◦ After these conditions are met, you can perform the “SVM Provision” operation, and then perform the “Export & Halt” operation to enable client access to ONTAP volumes. <p>b. Applicable for ONTAP releases 9.0 through 9.5:</p> <ul style="list-style-type: none"> ◦ Use the <code>vserver cifs modify</code> command to change the CIFS server name (CIFS Server NetBIOS Name). <p>Using this feature, you should create a CIFS server on the target SVM with a temporary identity, and then perform the “SVM Provision” operation.</p> <ul style="list-style-type: none"> ◦ After the “import” phase, you can run the <code>vserver cifs modify</code> command on the target cluster to replace the target SVM CIFS identity with the 7-Mode CIFS identity.

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

Supported and unsupported CIFS configurations for transition to ONTAP

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

Configurations that are supported for transition

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

- CIFS preferred DC configuration
- User mapping configuration:
 - `/etc/usermap.cfg`
 - `waf1.nt_admin_priv_map_to_root`
- CIFS local users and groups
- Symlink and widelink configuration (`/etc/symlink.translations`)
- CIFS audit configuration
- CIFS shares
- CIFS share ACLs
- CIFS home directory configuration
- CIFS options:
 - `cifs.gpo.enable`
 - `cifs.smb2.enable`
 - `cifs.smb2.signing.required`

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

See the precheck results for details about these CIFS configurations.

Configurations that are not supported in ONTAP

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

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

Machine accounts can access all shares after transition.

Configurations that must be manually transitioned

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

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

- Antivirus settings
- FPolicy configurations

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

- BranchCache configurations
- Character mapping configuration (charmap)
- Forcegroup attribute of CIFS shares to create files with a specified UNIX group as owning group
- Maxusers attribute of CIFS shares to specify the maximum number of simultaneous connections allowed to a 7-Mode CIFS share
- Storage-Level Access Guard (SLAG) configurations

- Share-level ACLs with UNIX-style permission
- Share ACLs for UNIX users and groups
- LAN Manager authentication level
- NetBIOS aliases
- CIFS search domains
- Some CIFS options

See the precheck results for details about these options.

Related information

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

Considerations for transitioning CIFS local users and groups

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

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

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

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

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

- A local user account with an empty password or local user accounts with passwords containing more than 14 characters on the 7-Mode system are transitioned to ONTAP software with the password **cifsUser@1**.

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

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

- If the 7-Mode Transition Tool IP address is not reachable from the target ONTAP software, the 7-Mode Transition Tool blocks the transition of CIFS local users and groups to the ONTAP software during the precheck phase. If you see this error during the precheck phase, use the

```
network ping -node local -destination ip_address
```

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

- The SVM to which the local users and groups are transitioned must have a data LIF.

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

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

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

Related information

[SMB/CIFS management](#)

Preparing for SAN transition

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

Creating SAN LIFs before transition

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

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

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

Steps

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

```
network interface create
```

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

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

```
network interface show
```

Related information

[SAN administration](#)

Configuring zones by using the FC zone plan

Before transitioning a SAN FC environment, you must configure zones by using the FC

zone planner to group the initiator hosts and targets.

- The FC zone planner must be generated by using the Collect and Access feature of the 7-Mode Transition Tool
- The FC zone script file must be accessible.
 1. If there are any changes to the igroup configurations on the 7-Mode systems, modify and regenerate the FC zone plan.

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

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

The following example runs the zone commands on the switch:

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

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

Preparing SAN hosts for transition

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

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

[Host and storage transition information collection](#)

Steps

1. Verify that the host is supported for transition.

[NetApp Interoperability Matrix Tool](#)

2. Perform the pretransition steps on the host.


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

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

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

Configurations that are transitioned

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

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

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

 - For ONTAP 8.3.2 and later supported releases, igroups and LUN mapping configurations are applied during the precutover operation.

Unsupported configurations in ONTAP

The unsupported configurations in ONTAP are as follows:

- 7-Mode Snapshot copy-backed LUN clones

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

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

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

- LUN clone split

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

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

- The `lun share` command
Sharing a LUN over NAS protocols
- SnapValidator

Configurations that must be manually transitioned

The following configurations must be transitioned manually:

- SAN LIFs
You must manually create the LIFs before transition.
- Portsets
You must manually configure igroups that are bound to a portset after transition.
- iSCSI access list information
- iSNS configuration
- iSCSI CHAP and RADIUS configurations

Related information

[NFS management](#)

[Network and LIF management](#)

Space considerations when transitioning SAN volumes

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

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

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

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

Related information

[Recovering from a failed LUN transition](#)

Preparing data protection features for transition

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

Preparing the cluster for transitioning volume SnapMirror relationships

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

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

Steps

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

```
system license add license_code
```

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

[Cluster and SVM peering express configuration](#)

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

```
job schedule cron create
```

Related information

[ONTAP 9 commands](#)

Data protection transition: supported and unsupported configurations

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

Supported configurations

You can transition volume SnapMirror relationships by using the 7-Mode Transition Tool. You can perform a copy-free transition of primary and secondary HA pairs. You must then manually set up the volume SnapMirror relationships after transition.

[Transitioning a SnapMirror relationship](#)

Unsupported configurations

- SnapVault relationships

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

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

- Qtree SnapMirror relationships

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

- Disaster recovery vFiler unit

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

- NDMP configuration

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

[Data protection using tape backup](#)

- Synchronous SnapMirror relationships


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

Related information

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

Transitioning 7-Mode aggregates using copy-free transition

The copy-free transition workflow involves planning a project, applying 7-Mode configurations to the SVMs, exporting the 7-Mode system information and halting the 7-Mode system, manually cabling the disk shelves to cluster nodes, and importing the 7-Mode data and configurations.



You must have prepared the 7-Mode system and the cluster for copy-free transition.

The storage cutover time can be 4-8 hours or less. The cutover time includes the time taken by the tool to perform two automated operations—the export and halt and import operations-- as well as the time taken for manually cabling the disk shelves to the new controllers.

The export and import operations take about 2 hours or less. Cabling can take 2-6 hours or less.

Planning a copy-free transition project

Planning a copy-free transition project involves selecting the source 7-Mode controllers and target cluster nodes, mapping 7-Mode volumes to a storage virtual machine (SVM), selecting the LIFs to be transitioned, and running prechecks.

You can create multiple projects with the same target cluster HA pair nodes. You can then run prechecks and apply the SVM configurations on all these projects. However, only one project can be in the critical section

window at a given time. A project is in the critical section window if the project is in any of the phases from export to commit, or if a rollback operation has been initiated for the project. You can proceed with the export and halt operation for another project only after the commit or rollback operation is completed for the project in the critical section window.

Copy-free transition project planning worksheets

You can use the copy-free transition planning worksheets to record information about node mapping, SVM mapping, volume mapping, and LIFs to transition. The worksheets are useful when creating a transition project by using the 7-Mode Transition Tool. You should be aware of the guidelines for completing the worksheets.

You can follow these guidelines to complete the worksheets:

- Map each vFiler unit to an SVM.

If there are no vFiler units in the 7-Mode controller, map the controller to a single SVM.
- Record the 7-Mode volume name and the corresponding Data ONTAP volume name.

The ONTAP volume name might be different from the 7-Mode volume name.

- Identify the LIFs to be configured on each SVM.

The IP addresses for the LIFs can either be existing on the 7-Mode system or can be new LIFs.

Node mapping

7-Mode controller	Mapped cluster node

SVM and volume mapping

7-Mode controller	vFiler unit or controller	Mapped SVM	7-Mode volume	SVM volume

LIF mapping (7-Mode IP addresses)

7-Mode controller	vFiler unit or controller	Mapped SVM	7-Mode IP address	Netmask	Default gateway	Home node	Home port

7-Mode controller	vFiler unit or controller	Mapped SVM	7-Mode IP address	Netmask	Default gateway	Home node	Home port

LIF mapping (new LIFs)

SVM	New IP address	Netmask	Default gateway	Home node	Home port

Example

The following table shows an example of completed worksheets for a 7-Mode HA pair with the controllers hostA_7mode and hostB_7mode.

Node mapping

7-Mode controller	Mapped cluster node
hostA_7mode	cluster1_01
hostB_7mode	cluster1_02

SVM and volume mapping

7-Mode controller	Mapped cluster node	vFiler unit or controller	Mapped SVM	7-Mode volume	SVM volume
hostA_7mode	cluster1_01	vfilerA	svm1	volA	volA
				volB	volB
		vfilerB	svm2	vol1	vol_nfs
				vol2	vol_cifs
hostB_7mode	cluster1_02	Not applicable	svm3	vol3	vol3
				vol4	vol4
				vol5	vol5
				vol6	vol6

LIF mapping (7-Mode IP addresses)

7-Mode controller	vFiler unit or controller	Mapped SVM	Existing 7-Mode IP	Netmask	Default gateway	Home node	Home port
hostA_7mode	vfilerA	svm1	192.0.2.129	255.255.255.128	192.40.0.1	cluster1_01	e0a
			192.0.2.135	255.255.255.128	192.40.0.1	cluster1_02	e0b
	vfilerB	svm2	-				
			-				
hostB_7mode	Not applicable	svm3	192.0.2.110	255.255.255.128	192.40.0.1	cluster1_01	e0c
			192.0.2.111	255.255.255.128	192.40.0.1	cluster1_02	e0d

LIF mapping (new LIFs)

SVM	New IP address	Netmask	Default gateway	Home node	Home port
svm1	-				
	-				
svm2	192.0.2.130	255.255.255.128	192.40.0.1	cluster1_01	e1c
	192.0.2.131	255.255.255.128	192.40.0.1	cluster1_02	e1d
svm3	192.0.2.136	255.255.255.128	192.40.0.1	cluster1_01	e0c
	192.0.2.137	255.255.255.128	192.40.0.1	cluster1_02	e0d

Adding controllers and clusters

Before you start the transition, you must add the 7-Mode controllers, including both nodes of a 7-Mode HA pair, and the clusters that are required for the transition. You should add the clusters using the cluster-management interface.

- For a copy-free transition, you must add the cluster, and not the cluster nodes that are the target of transition.
- The 7-Mode controllers and clusters information that you provide is not persistent.

If the 7-Mode Transition Tool service is restarted, the tool prompts in the project dashboard for information about controllers and cluster that are part of active projects.


Steps

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

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

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

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

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

The Select Source Systems screen is displayed.

Creating a copy-free transition project

The first step in planning a transition project is to select the source 7-Mode HA pair from which you want to transition the disk shelves, aggregates, volumes, and configurations, and then create a transition project.

- The 7-Mode controllers in the HA pair must be running a supported ONTAP version on a platform that is supported for a copy-free transition.

[NetApp Interoperability Matrix Tool](#)

- Both controllers in the HA configuration must be healthy.
 1. Select the **Copy-Free Transition** migration method from the homepage and click **Start Planning**.

If the controller and cluster required for a new project are not added, you can enter the details in the Enter Device Credentials pane.

2. Select the source 7-Mode HA pair that you want to transition.
3. Click **Create Project**.
 - a. In the Project Details window, provide a name for the project.
 - b. Select a project group to which the project should be added.

You can either create a new project group or add the project to the default group.

Creating a project group enables you to group and monitor related projects.

- c. Click **Save**.

The Select Target Cluster screen is displayed.

Selecting the target cluster nodes for transition

You can select the target cluster HA pair and map each 7-Mode controller in the HA pair to a corresponding target cluster node. The mapped node specifies the cluster node to which the disk shelves from the corresponding 7-Mode controller must be connected.

The target cluster must be running Data ONTAP 8.3.2 or later.

You can transition the 7-Mode disk shelves to a target HA pair that has preexisting data aggregates and volumes.

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

Steps

1. Select the target HA pair to which the 7-Mode disk shelves must be connected.

The tool automatically maps each 7-Mode storage system to a target cluster node.



The disk and aggregate ownership from each 7-Mode controller is transferred to its corresponding mapped target cluster node during the import phase.

2. Click **Swap Node Mapping** to change the automatic assignment of source-to-target node mapping.
3. Click **Save and Continue**.

The SVM and Volume Mapping screen is displayed.

Mapping SVMs and volumes

You should map each 7-Mode controller in the HA pair to a target SVM. If you have vFiler units, you should select a target SVM for each vFiler unit. The volumes from the 7-Mode controller or vFiler unit are transitioned to the mapped SVM.

You must have created the SVMs on the target cluster.

Cluster management with System Manager

System administration

A vFiler unit can be mapped only to a single SVM. Volumes from any other 7-Mode controller or vFiler unit cannot be transitioned to a mapped SVM.

Steps

1. In the SVM and Volume Mapping tab, select the target SVM to which you want to transition the volumes from each 7-Mode controller or vFiler unit.

The target SVM can be in the default or non-default IPspace.

2. Depending on whether you want to apply the same junction path policy for the volumes in all the SVMs or a different junction path policy for the volumes in each SVM, choose one of the following actions:

If you want to...	Then...
Apply the same junction path policy to all the SVMs	a. Click Apply . b. Select an option for the junction path policy. c. Click OK .
Specify the junction path policy for each SVM	Select the junction path policy from the drop-down list for each target SVM.

The junction path policy specifies the path with which the target clustered Data ONTAP volumes must be mounted for client access. You can add one of the following values for the junction path policy:

- **Preserve 7-Mode mount paths**

Retains the same junction paths as that being used on the source 7-Mode volumes and the volumes are mounted with junction paths in the format */vol/source_volume_name* after transition.

- **Use clustered Data ONTAP volume name**

All the target clustered Data ONTAP volumes are mounted with junction paths with the clustered Data ONTAP volume name in the format */target_volume_name* after transition.

- **Use 7-Mode volume name**

All the target clustered Data ONTAP volumes are mounted with junction paths with the 7-Mode volume name in the format */source_volume_name* after transition.

3. Click  to modify the name of the target clustered Data ONTAP volume.

By default, the target clustered Data ONTAP volume has the same name as the 7-Mode volume. If a volume with the same name as the 7-Mode volume already exists on the SVM, the target volume is automatically assigned a new name.

4. Click **Save Mapping** for each mapped SVM.
5. Click **Next**.

The Networking screen is displayed.

Selecting LIFs for transition

You can optionally specify the LIFs that you want to configure on the SVMs after transition. These LIFs can be existing IP addresses on the 7-Mode systems or they can be new LIFs. Only NAS LIFs are transitioned. FC and iSCSI LIFs must be manually configured before the SVM provision phase.

The LIFs that are selected for transition are configured on the SVMs during the SVM provision phase in the following ways:

- Existing 7-Mode IP addresses that are selected for transition are created in the administrative down state.

These IP addresses can continue to serve data in 7-Mode until the cutover starts. During the import phase, these IP addresses are configured in the administrative up state.

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

You can use these LIFs to test the connectivity of the SVMs to the name servers after the SVM provision phase.

Steps

1. In the LIF configuration tab, choose one of the following options:

If you want to transition...	Then...
An existing IP address on the 7-Mode system	<ol style="list-style-type: none">a. Click Select 7-Mode LIF.b. Select the IP address that you want to transition, and then specify the target SVM and other network parameters.c. Click Save.
A new IP address	<ol style="list-style-type: none">a. Click Add New LIF.b. Specify the IP address that you want to configure, the target SVM, and other network parameters.c. Click Save.



The target ports must be in the same IPspace as the target SVM.

2. Click **Next**.

The Plan Configuration tab is displayed.

Customizing the transition of 7-Mode configurations

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

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

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

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

Steps

1. From the Plan Configuration page, select the following options from the **SVM Configuration** pane:

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

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

Related information

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

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

[Data protection transition: supported and unsupported configurations](#)

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

Examples of consolidating NFS export rules and Snapshot schedules for transition

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

Example of consolidating NFS export rules for transition

NFS export rules in 7-Mode and ONTAP before transition

7-Mode export rules

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

Export policies existing in ONTAP

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

The existing export policy export_policy_1 has the following export rule:

```
cluster-2::> vserver export-policy rule show -vserver vs1 -policyname
export_policy_1
Policy      Rule    Access  Client      RO
Vserver     Name      Index   Protocol  Match      Rule
-----
vs1         export_policy_1 1      nfs      0.0.0.0/0   sys
```

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

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

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

```
cluster-2::> vserver export-policy rule show -vserver vs1 -policyname
transition_export_policy_1
Policy      Rule    Access  Client      RO
Vserver     Name      Index   Protocol  Match      Rule
-----
vs1         transition_export_policy_1
                        1      nfs      0.0.0.0/0   sys
```

```
cluster-2::> volume show -vserver vs1 -volume vol1,vol2,vol3 -fields
policy
vserver volume policy
-----
vs1      vol1      transition_export_policy_1
vs1      vol2      transition_export_policy_1
vs1      vol3      transition_export_policy_1
3 entries were displayed.
```

Export policies in ONTAP after transition with consolidation and reuse

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

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

```
cluster-2::> vserver export-policy rule show -vserver vs1 -policyname
export_policy_1
Vserver          Policy          Rule      Access   Client      RO
Vserver          Name              Index     Protocol Match      Rule
-----
vs1              export_policy_1 1         nfs       0.0.0.0/0    sys
```

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

Example of consolidating Snapshot policies for transition

Snapshot schedules in 7-Mode and ONTAP before transition

7-Mode schedule

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

Snapshot policies existing in ONTAP

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

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

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

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

Snapshot policy in ONTAP after transition with consolidation and reuse

7-Mode volume	7-Mode Snapshot schedule	Snapshot policy in ONTAP
vol1	0 2 4@8,12,16,20 (weekly Snapshot copies: 0, daily Snapshot copies: 2, hourly Snapshot copies: 4 at 2, 4, 8, 12, 16, 20 hours)	Consolidated policy for vol1, vol2, and vol3 for which the existing ONTAP policy is reused Name: ScheduleDailyHourly4
vol2	0 2 4@8,12,16,20	vol3
0 2 4@8,12,16,20	vol4	1 2 3@8,12,16 (weekly Snapshot copies: 1, daily Snapshot copies: 2, hourly Snapshot copies: 3 at 8,12,16 hours)
<ul style="list-style-type: none"> • Name: transition_snapshot_policy_1 • Schedule details <ul style="list-style-type: none"> ◦ Schedule1: weekly, count1: 1 ◦ Schedule2: daily, count2: 2 ◦ Schedule3: hourly, count3: 3 every 8,12,16 hours 	vol5	2 2 3@8,12,16 (weekly Snapshot copies: 2, daily Snapshot copies: 2, hourly Snapshot copies: 3 at 8,12,16 hours)

Running prechecks

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

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

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

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



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

Steps

1. 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**.
2. To save the raw results in comma-separated values (CSV) format and export the results, click **Save As CSV**.

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

You must resolve all the errors detected by the prechecks before you start data copy. It is also a best practice to resolve all warnings prior to proceeding with the migration process. Resolution can be resolving the source issue of the warning message, implementing a workaround, or accepting the result of the issue.

Severity levels for precheck messages

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

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

- **Error**

Configurations that cannot be transitioned.

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

- Traditional volumes
- SnapLock volumes
- Offline volumes

- **Warning**

Configurations that can cause minor problems after transition.

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

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

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

- **Information**

Configurations that have been successfully transitioned.

Applying SVM configurations

You must manually apply some configurations before the SVM provision phase. You can then apply all configurations that are defined in the 7-Mode controller (files in the `/etc` directory) or at the vFiler unit level to the mapped SVMs by using the tool.

If you do not want to transition all configurations to the target SVMs, you must have set the properties for the configurations to be excluded on the SVMs.

Customizing the transition of 7-Mode configurations

- Configurations such as NFS exports, CIFS shares, and LUN mapping are not applied to the SVM in the SVM provision phase.
- The following configurations are applied by the tool on the SVMs in the SVM provision phase:
 - **Name services**
 - DNS configuration
 - LDAP configuration
 - NIS configuration
 - Name service switch configuration
 - Hosts configuration
 - UNIX users and groups
 - Netgroups configuration

- **Networking**

- Existing 7-Mode IP addresses that are selected for transition are created in the administrative down state.

During the import phase, these IP addresses are configured in the administrative up state.

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

- **NFS**

NFS options

- **CIFS**

- CIFS preferred DC configuration
- User mapping configuration
- Widelinks configuration
- CIFS options
- Audit configuration

- **SAN**

FC and iSCSI services



You cannot rerun this operation after it is completed successfully.

Therefore, if you make any changes to the 7-Mode controller-level configurations after this operation, you must manually transition the configurations to the target SVMs before the export phase. For example, if you add a new vFiler unit to the 7-Mode controller and map it to an SVM after this operation, you must manually transition the configurations of that vFiler unit to the mapped SVM. Another example is if you add some UNIX users on the 7-Mode controller after the SVM provision phase, you must create these UNIX users manually on the mapped SVMs.

Steps

1. Click **Apply SVM Config** to apply the 7-Mode configurations to the target SVMs.

A confirmation message that lists important considerations for this operation is displayed.

2. Click **Yes** to continue.
3. After the operation is complete, perform the following steps:
 - a. Click **Save as CSV** to save the operation results in a file.
 - b. Click **Collect Project Logs** to create a backup of all of the transition log files.

It is a best practice to save the log files after each transition operation.

- c. Click **Close** to close the operation results window.

If the operation takes a long time to complete, you can click **Run in Background** to exit the operation results window. You should not edit the project or perform any other task when the operation is running in the background. You can then view the operations results from the Operation History tab.

4. Verify and test the configurations that are applied to the target SVMs manually and make the required changes.
5. Manually verify the connectivity to external name servers.

Verifying that 7-Mode systems are ready for cutover

Before disconnecting client access, you can verify the readiness of the 7-Mode system for storage cutover, such as verifying whether the SP or RLM is configured on the 7-Mode system and whether the disk shelves, aggregates, and volumes are ready for transition. You can manually fix any issues before the export operation, thereby reducing the downtime.

You can run this operation multiple times before the export and halt operation is initiated.

Steps

1. Click **Check Readiness** to verify that the 7-Mode systems are ready for cutover.

Although this operation is optional, it is a best practice to verify the readiness of the systems and fix issues before the export operation in order to minimize the storage cutover window.

2. Wait for the operation to complete and do the following:

- a. Click **Save as CSV** to save the operation results in a file.
- b. Click **Collect Tool Logs** to take a backup of all the transition log files.

It is a best practice to save the log files after each transition operation.

- c. Click **Close** to close the operation results window.

If the operation takes a long time to complete, you can click **Run in Background** to exit the operation results window. You should not edit the project or perform any other task when the operation is running in the background. You can then view the operations results from the Operation History tab.

Exporting storage configurations and halting 7-Mode systems

The cutover window for the transition starts from the export phase. In this phase, the tool collects system information, disk shelf details, and storage configurations from the 7-Mode systems, and then halts the 7-Mode storage systems.

- The Service Processor (SP) or Remote LAN Module (RLM) must be configured with an IPv4 address on the 7-Mode system.
- All clients must be disconnected from the 7-Mode systems (by unmounting NFS exports, disconnecting CIFS shares, and shutting down SAN hosts), but the applicable NAS and SAN services must be running on the 7-Mode systems.



You must not stop any protocol services because they are required for collecting the protocol configurations from the 7-Mode storage systems.

- Any data written to the 7-Mode volumes during this operation is lost.
- You must not perform any management operations on the 7-Mode systems during this operation.

- The tool performs the following operations in the export phase:

- Collects all volume and storage configurations
- Creates a Snapshot copy of each transitioning aggregate

If you decide to revert to 7-Mode at any stage of transition, these Snapshot copies are used by the tool for the revert operation.

- Boots the 7-Mode controllers in maintenance mode
- Removes disk ownerships from the 7-Mode controllers
- Disables disk autoassignment on the target cluster nodes

Steps

1. Click **Export & Halt**.

A message that lists important considerations for this operation is displayed.

2. Click the **Confirm that the client access is disconnected** check box.
3. Click **Yes** to continue with the export and halt operation.

The operation results are displayed.

4. Wait for the operation to complete and perform the following steps to save the operation results and collect the tool logs:
 - a. Click **Save as CSV** to save the operation results in a file.
 - b. Click **Collect Tool Logs** to create a backup of all the transition log files.

It is a best practice to save the log files after each transition operation.

- c. Click **Close** to close the operation results window.

If the operation takes a long time to complete, you can click **Run in Background** to exit the operation results window. You should not edit the project or perform any other task when the operation is running in the background. You can then view the operations results from the Operation History tab.

Disconnecting disk shelves from the 7-Mode system and connecting to cluster nodes

Cabling the 7-Mode disk shelves to the target cluster nodes is a manual process. After cabling the disk shelves, it is a best practice to verify the cabling by using Config Advisor. You can then verify the cabling by using the 7-Mode Transition Tool. The tool performs only a subset of the checks that are performed by Config Advisor.

You must have recorded the information about disk shelf connectivity to the 7-Mode controller ports.


You must be aware of some of the considerations for connecting SAS disk shelves:

- You must follow the rules for cabling SAS square and circle ports.
- IOM6 and IOM3 shelves can be mixed in the same stack, but there should be no more than a single transition between shelves that are using different IOM types.

For example, IOM6e (controller)--IOM6 (shelf)--IOM3 (shelf)--IOM3 (shelf) is a supported configuration. But IOM6e (controller)--IOM3 (shelf)--IOM6 (shelf)--IOM3 (shelf) is not a supported configuration.

Steps

1. Check the disk shelf IDs in the 7-Mode HA pair and the target cluster nodes:
 - a. If there are duplicate shelf IDs (if the 7-Mode disk shelf IDs are used for the disk shelves in the target cluster nodes), change the disk shelf IDs.
 - For SAS disk shelves, a valid shelf ID is 00 through 99.
 - SAS shelf IDs must be unique within the HA pair. [SAS Disk Shelves Installation and Service Guide for DS4243, DS2246, DS4486, and DS4246](#)
 - For FC disk shelves, a valid shelf ID is 1 through 7.
 - FC shelf IDs must be unique within each FC loop. [DS14mk2 FC, and DS14mk4 FC Hardware Service Guide](#)
 - b. Power cycle the disk shelves for the new IDs to take effect.
2. Power off the 7-Mode disk shelves.
3. Depending on whether additional ports are available on the target cluster nodes, choose one of the following options:

If...	Then...
Additional ports are available to connect the disk shelves	Connect the disk shelves in a new stack to the target cluster nodes in a multipath configuration. <div>  <p>It is the best practice to connect the disk shelves to a separate stack. The 7-Mode Transition Tool detects the availability of the required number of ports on the target cluster nodes when prechecks are run.</p> </div>
Ports are not available to connect the disk shelves	Perform one of the following actions: <ul style="list-style-type: none"> • Add a new expansion card and connect the disk shelves in a new stack to the target cluster nodes in a multipath configuration. <p>You must have verified that the expansion card is supported on the destination platform.</p> <ul style="list-style-type: none"> • Connect the disk shelves to an existing stack in a multipath configuration.

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

[DiskShelf14mk2 AT Hardware Service Guide](#)

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

4. Power on the disk shelves.



You must wait for at least 70 seconds before you proceed.

5. Use Config Advisor to verify the connections.

You must fix any cabling issues identified by Config Advisor.

[NetApp Downloads: Config Advisor](#)

6. From the 7-Mode Transition Tool, click **Verify Cabling**.

A message listing the important considerations for this operation is displayed.

Troubleshooting: If there are missing disks in an aggregate, the aggregate becomes degraded and the cabling verification fails. If the number of missing disks is within a permissible limit, you can continue the transition with the degraded aggregates by running the following command from the 7-Mode Transition Tool CLI:

```
transition cft aggregate degraded-transition -p project_name -n 7-  
mode_host_name -a 7-mode_aggregate_name -i acknowledge
```

You can then rerun the cabling verification operation and continue with the transition. You must ensure that there are enough spare disks in the target cluster node to reconstruct these RAID groups after the aggregates are transitioned.

7. Click **Yes** to continue.
8. Wait for the operation to complete and perform the following steps to save the operation results and collect the tool logs:
 - a. Click **Save as CSV** to save the operation results in a file.
 - b. Click **Collect Tool Logs** to create a backup of all the transition log files.

It is a best practice to save the log files after each transition operation.

- c. Click **Close**.

If the operation takes a long time to complete, you can click **Run in Background** to exit the operation results window. You should not edit the project or perform any other task when the operation is running in the background. You can then view the operations results from the Operation History tab.

Related information

[Gathering cabling information for transition](#)

Importing 7-Mode data to ONTAP

After verifying the cabling and resolving any issues, you can run the import operation. In this phase, the disk ownership is assigned to the mapped cluster nodes, and the 7-Mode aggregates, volumes, and LUNs are converted to the ONTAP format. All the volume-level and LUN-level configurations are also applied.

The following operations are performed by the tool in this phase:

- 7-Mode disks are assigned to the mapped target cluster nodes.
- All the 7-Mode aggregates, volumes, and LUNs are converted to the ONTAP format.
- The 7-Mode IP addresses that were selected for transition are configured on the SVMs in the administrative up state.
- The following configurations are applied:
 - NFS export rules
 - CIFS shares
 - CIFS ACLs configuration
 - CIFS home directory configuration
 - CIFS symbolic links
 - Quota configuration
 - Snapshot copy schedules
 - LUN maps and igroups

Steps

1. Click **Import**.

Troubleshooting: If there are missing disks in an aggregate, the aggregate becomes degraded and the import operation fails. If the number of missing disks is within a permissible limit, you can continue the transition with the degraded aggregates by running the following command from the 7-Mode Transition Tool CLI:

```
transition cft aggregate degraded-transition -p project-name -n 7-mode-host-name -a 7-mode-aggregate-name -i acknowledge
```

You can then rerun the cabling verification operation and continue with the transition. You must make sure that there are enough spare disks in the target cluster node to reconstruct these RAID groups after the aggregates are transitioned.

A warning message is displayed, listing the important considerations for this operation.

2. Click **Yes** to continue.

The operation results are displayed.

3. Wait for the operation to complete and do the following:

- a. Click **Save as CSV** to save the operation results in a file.
- b. Click **Collect Tool Logs** to take a backup of all the transition log files.

It is a best practice to save the log files after each transition operation.

- c. Click **Close** to close the operation results window.

If the operation takes a long time to complete, you can click **Run in Background** to exit the operation results window. You should not edit the project or perform any other task when the operation is running in the background. You can then view the operations results from the Operation History tab.


Completing the transition


Completing the transition involves manually verifying the transitioned volumes and configurations, testing your workload, starting production, and then committing the copy-free transition project. Because rollback is not allowed after the commit operation, you should verify all workloads and start production for a brief duration to evaluate if rollback is required.

Restrictions during preproduction testing

Some operations are blocked and some operations are not recommended during preproduction testing. These restrictions are imposed for allowing a rollback to 7-Mode if you do not want to commit the transition.

Operations that are blocked

Operation	Description
Automatically deleting (autodelete) aggregate Snapshot copies	<p>Because the 7-Mode aggregate-level Snapshot copies created during the export operation are required in case of a rollback, the Snapshot copies are not deleted automatically when the used space in the aggregate grows.</p> <div><p>You must monitor the free physical space in the aggregate and ensure that the aggregates do not run out of space during the testing.</p></div>
Moving volumes to another aggregate	<ul style="list-style-type: none">• You cannot move volumes to the transitioned aggregates.• You can move volumes from the transitioned aggregates to aggregates in the cluster.
Copying or moving LUNs across volumes	<ul style="list-style-type: none">• You cannot copy or move LUNs to the transitioned volumes.• You can copy or move LUNs from the transitioned volumes to other volumes in the cluster.
Creating an aggregate	<p>This operation is restricted on the target cluster HA pair. You can create aggregates on other nodes in the cluster.</p>
Destroying an aggregate	<p>Because the 7-Mode aggregate-level Snapshot copies created during the export operation are required in case of a rollback, the transitioned aggregates cannot be destroyed.</p>

Operation	Description
Setting up a transitioned aggregate as the root aggregate	The transitioned aggregates cannot be selected as root aggregates. Additionally, you cannot modify the HA policy of the transitioned aggregates to CFO.
Performing file copy operations	<ul style="list-style-type: none"> You cannot move or copy files (single file copy-on-demand) to the transitioned volumes. You can move or copy files from the transitioned volumes to other volumes in the cluster.
Mirroring an existing aggregate	The operation is blocked on all the aggregates in the cluster.
Upgrading or reverting the Data ONTAP version on the target cluster nodes	You must commit the project before upgrading or reverting the target cluster nodes.
Adding disks	<p>You cannot run the <code>storage aggregate add-disks</code> command in the admin privilege level. However, you can run this command in the advanced privilege level.</p> <div>  <p>You must ensure that only the spare disks from the 7-Mode disk shelves are added for increasing the space in the transitioned aggregates. You must add spare disks by using the <code>-disklist</code> parameter (the <code>-diskcount</code> parameter must not be used).</p> </div>
Designating a transitioned volume as the SVM root volume	You cannot run the <code>volume make-vsroot</code> command on the transitioned volumes.

Operations that are not recommended

Operation	Corrective action before rollback
<p>Relocating aggregates</p> <p>The ownership of the transitioned aggregate is changed to its HA partner.</p>	Swap the aggregate ownership before rollback because the 7-Mode Transition Tool maps the aggregates to the 7-Mode controllers based on the target node mapping information saved in the project.
Creating volumes on the transitioned aggregates	You must delete these volumes or move them to different aggregates.
Renaming aggregates or volumes	Rename the aggregates or volumes to their original names.

Operation	Corrective action before rollback
Changing the RAID type	The RAID type must match the 7-Mode RAID type if you decide to roll back to 7-Mode.

Related information

[ONTAP 9 commands](#)

Rehosting transitioned volumes to a different SVM

Volume rehost enables you to migrate a transitioned volume from one SVM to another SVM without data copy. The rehost operation enables you to consolidate all volumes that have FC LUNs to a single SVM, thereby preserving the 7-Mode single-system image (SSI) semantics. You can also rehost transitioned NAS volumes.

- The volume that you want to rehost must be online.
- Volume management operations, such as volume move or LUN move, must not be running.
- Data access to the volume that is being rehosted must be stopped.

Rehosting is a disruptive operation.

The following volume policies, policy rules, and configurations are lost from the source volume and must be manually reconfigured on the rehosted volume after the rehost operation:

- Volume and qtree export policies
- Antivirus policies
- Volume efficiency policy
- Quality of Service (QoS) policies
- Snapshot policies
- Quota rules
- CIFS shares
- igroups associated with a portset

Steps

- **Rehosting FC and iSCSI volumes**
 - a. Switch to the advanced privilege level:
set -privilege advanced
 - b. Rehost the volume on the destination SVM:

If you want to...	Run the following command...
Unmap the LUNs after rehosting	volume rehost -vserver <i>source_svm</i> -volume <i>vol_name</i> -destination-vserver <i>destination_svm</i> -force-unmap-luns true
Remap the LUNs to the same igroups after rehosting	volume rehost -vserver <i>source_svm</i> -volume <i>vol_name</i> -destination-vserver <i>destination_svm</i> -auto-remap-luns true

• Rehosting NFS volumes

- Record information about the NFS export policies.
- Unmount the volume from the parent volume:
volume unmount
- Switch to the advanced privilege level: **set -privilege advanced**
- Rehost the volume on the destination SVM:
volume rehost -vserver *source_svm* -volume *vol_name* -destination-vserver *destination_svm*

The default export policy of the destination SVM is applied to the rehosted volume.

- Create the export policy:
vserver export-policy create
- Update the export policy of the rehosted volume to a user-defined export policy:
volume modify
- Mount the volume under the appropriate junction path in the destination SVM:
volume mount
- Verify that the NFS service is running on the destination SVM:
vserver nfs status
- Resume NFS access to the rehosted volume.

Because the volume access path (LIFs and junction path) has undergone changes, you must update the NFS client credentials and LIF configurations to reflect the destination SVM LIFs.

NFS management

• Rehosting CIFS volumes

- Record information about the CIFS shares.
- Unmount the volume from the parent volume:
volume unmount
- Switch to the advanced privilege level:
set -privilege advanced
- Rehost the volume on the destination SVM:

```
volume rehost -vserver source_svm -volume vol_name -destination-vserver destination_svm
```

- e. Mount the volume under the appropriate junction path in the destination SVM:

```
volume mount
```

- f. Create CIFS shares for the rehosted volume:

```
vserver cifs share create
```

- g. If the DNS domains differ between the source and destination SVMs, create new users and groups.

- h. Update the CIFS client with the new destination SVM LIFs and junction path to the rehosted volume.

[SMB/CIFS management](#)

• Rehosting volumes in SnapMirror relationships

- a. Record the SnapMirror relationship type:

```
snapmirror show
```

- b. From the destination cluster, delete the SnapMirror relationship:

```
snapmirror delete
```



You must not break the SnapMirror relationship; otherwise, the data protection capability of the destination volume is lost and the relationship cannot be reestablished after the rehosting operation.

- c. From the source cluster, release the SnapMirror relationship information:

```
snapmirror release
```

You set the `-relationship-info-only` parameter to true so that the Snapshot copies are not deleted and only the source relationship information is removed.

- d. Switch to the advanced privilege level:

```
set -privilege advanced
```

- e. Rehost the volume on the destination SVM:

```
volume rehost -vserver source_svm -volume vol_name -destination-vserver destination_svm
```

- f. Create the SVM peer relationship between the source and destination SVMs:

```
vserver peer create
```

- g. Create the SnapMirror relationship between the source and destination volumes:

```
snapmirror create
```

The rehosted volume can be the source or destination of the SnapMirror relationship.

- h. Resynchronize the data protection relationship:

```
snapmirror resync
```

[Data protection](#)

You must manually create the autovolume workloads for the rehosted volumes by performing the following steps:

1. Create a user-defined policy group for the SVM:

```
qos policy-group create -vserver destination-vserver -policy-group policy-group-name
```

2. Assign the QoS policy group to the rehosted volume:

```
volume modify -vserver destination-vserver -volume rehosted-volume -qos-policy-group policy-group-name
```

You must manually reconfigure the policies and the associated rules on the rehosted volume.



If the rehosting operation fails, you might need to reconfigure the volume policies and the associated rules on the source volume.

Related information

[ONTAP 9 commands](#)

Verifying the transitioned configurations

After successfully importing the 7-Mode volumes and configurations, you must manually verify the transitioned aggregates, volumes, LUNs, and configurations.

Steps

1. Verify that the 7-Mode aggregates, volumes, and LUNs as well as the CIFS shares, NFS exports, and LUN mappings are transitioned.
2. Verify that all the 7-Mode configurations are retained.

Performing manual configuration tasks after transition

You must manually perform some configuration tasks that are required for the workloads and applications that are accessing the transitioned volumes. You can obtain the list of manual tasks from the precheck results.

Steps

1. Perform the tasks listed in the precheck results for configuring features that were not transitioned by the tool or that require customization for your environment.

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

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

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

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

[Transitioning a SnapMirror relationship](#)

Testing the workloads and applications

You should manually test all workloads and applications in the preproduction

environment. You can then start production for a brief duration to evaluate if rollback is required before committing the project.

The transitioned aggregates must have at least 5% free physical space.



The best practice is to have at least 20% free space in the transitioned aggregates.

Some operations are restricted during preproduction testing.

Restrictions during preproduction testing

Steps

1. Connect clients to the transitioned volumes.
2. If you have SAN workloads, perform the post-transition host remediation tasks on the SAN hosts.

SAN host transition and remediation

3. Test all the workloads and applications that use the transitioned data and configurations.
4. Verify that the transitioned aggregates are not running out of space by monitoring the free physical space in the transitioned aggregates from the Aggregates tab in the 7-Mode Transition Tool dashboard.

Troubleshooting: If you run out of space on the transitioned aggregates, you can add disks.

- a. Log in to the advanced privilege level:

```
set -privilege advanced
```

- b. Select the spare disks from the 7-Mode disk shelves and add disks to increase the space in the transitioned aggregates:

```
storage aggregate add-disks -aggregate aggr_name -disklist disk1
```

If 7-Mode spare disks are not available, you can use spare disks from the disk shelves in the cluster nodes; however, doing so complicates the rollback process.

You can start serving production data.



You can serve data in the production environment for a brief duration to ensure that the workloads are operating correctly in a production environment and a rollback to 7-Mode is not required. You must not prolong this phase and must not delay committing the copy-free transition project for the following reasons:

- The probability of running out of space in the transitioned aggregates increases as new data is written to the volumes.
- Any new data written to the volumes during this stage will not be available during rollback.

Related information

[Performing a transition rollback to 7-Mode](#)

[ONTAP 9 commands](#)

Committing the copy-free transition project

The final step in transition is to commit the copy-free transition project. After committing the aggregates, you cannot perform a rollback to 7-Mode.

You must have manually verified the transitioned data and configurations and tested workloads and applications.

All the aggregate-level Snapshot copies that were created in the export phase are deleted.

Steps

1. Click **Commit**.
2. In the warning message that is displayed, click **Yes**.

All the preproduction testing phase restrictions are removed and the transitioned volumes can serve production data, if not done in the preproduction testing phase.

Transitioning a SnapMirror relationship

You can transition the secondary HA pair first, set up a staggered SnapMirror relationship between the 7-Mode primary volumes and clustered Data ONTAP secondary volumes, and then transition the primary HA pair later. In some scenarios, you must transition the secondary and primary HA pairs of a 7-Mode SnapMirror relationship in parallel.

If all the primary volumes belong to one HA pair and all the secondary volumes belong to the other HA pair, you can transition by using the staggered method.

If either HA pair has a mix of primary and secondary volumes, you must transition by using the parallel method.

After transitioning the primary and secondary HA pairs, you must manually set up the volume SnapMirror relationship in clustered Data ONTAP after transition. For a successful resynchronization, at least one common Snapshot copy, which is created in Data ONTAP 8.1 or later, must exist between the primary and secondary volumes of the SnapMirror relationship.

Related information

[Cluster management with System Manager](#)

Transitioning HA pairs in a SnapMirror relationship in a staggered configuration

You can transition the secondary HA pair first, set up a staggered SnapMirror relationship between the 7-Mode primary volumes and ONTAP secondary volumes, and then transition the primary HA pair later.

You must have prepared the source and destination clusters for transitioning the SnapMirror relationships.

[Preparing the cluster for transitioning volume SnapMirror relationships](#)

Steps

1. From the 7-Mode Transition Tool, perform a copy-free transition of the HA pair that contain the secondary volumes of the 7-Mode volume SnapMirror relationships.

Before transitioning the 7-Mode HA pair that contains the secondary volumes, no manual intervention is required for the 7-Mode SnapMirror relationships. This ensures that the 7-Mode secondary volumes are transitioned as read-only volumes to ONTAP.

Transitioning 7-Mode aggregates using copy-free transition

2. During the precommit testing phase of the secondary HA pair, create a disaster recovery relationship between the 7-Mode primary volume and ONTAP secondary volume:
 - a. From the secondary destination cluster, use the `vserver peer transition create` command to create an SVM peer relationship between the 7-Mode primary volume and the ONTAP secondary volume.
 - b. Use the `job schedule cron create` command to create a job schedule that matches the schedule configured for the 7-Mode SnapMirror relationship.
 - c. Use the `snapmirror create` command to create a SnapMirror relationship of type TDP between the 7-Mode primary volume and the ONTAP secondary volume.
 - d. Use the `snapmirror resync` command to resynchronize the ONTAP secondary volume.

For successful resynchronization, a common 7-Mode Snapshot copy must exist between the 7-Mode primary volume and the ONTAP secondary volume.

3. Perform the required testing of the transitioned aggregates and volumes.
4. From the 7-Mode Transition Tool, commit the transition of the project for the secondary HA pair.

Committing the copy-free transition project

5. Perform a copy-free transition of the HA pair that contains the primary volumes of the 7-Mode volume SnapMirror relationships.

Transitioning 7-Mode aggregates using copy-free transition

6. During the precommit testing phase of the primary HA pair, create a SnapMirror relationship between the transitioned secondary and primary volumes.
 - a. From the destination cluster, create an intercluster SVM peer relationship between the SVMs that contain the transitioned primary and secondary volumes.

System administration

- b. Use the `snapmirror delete` command to delete the TDP SnapMirror relationship between the 7-Mode primary volume and the ONTAP secondary volume that was created in Step [#SUBSTEP_D528769DF8EC49058D1958565914CF47](#).
- c. Delete the cron job schedule that was created in Step [#SUBSTEP_EB470706425C45759EAAE8F0A87BA547](#):
job schedule cron delete
- d. Create a volume SnapMirror relationship between the transitioned primary and secondary volumes.

Volume disaster recovery express preparation

- e. On the destination volume, resynchronize the source volume and destination volume of the SnapMirror relationship:
snapmirror resync



At least one common Snapshot copy must exist between the source and destination volumes.

Troubleshooting: SnapMirror resynchronization fails if the common Snapshot copy was created in a Data ONTAP release earlier than 8.1. You can use the `-fs-version` parameter with the `volume snapshot show` command in the advanced privilege level to view the release in which the Snapshot copy is created. If this issue occurs, break the SnapMirror relationship and then perform the resynchronization.

- f. Monitor the status of the SnapMirror data transfers:

snapmirror show



You must not perform any operation, such as volume move or SnapMirror break, on the source and destination volumes until the resynchronization is completed successfully. Ensure that the resynchronization does not get aborted and completes successfully; otherwise, the volumes can go to an inconsistent state.

7. Commit the transition of the primary project.

[Committing the copy-free transition project](#)

Related information

[7-Mode data transition using SnapMirror](#)

[ONTAP 9 commands](#)

Transitioning primary and secondary systems in a SnapMirror relationship in parallel

You can transition the primary and secondary HA pairs that contain the volumes of a 7-Mode SnapMirror relationship in parallel. You must then manually set up the volume SnapMirror relationship in clustered Data ONTAP after transition. The SnapMirror relationship is retained after transition without requiring a rebaseline.

You must have prepared the source and destination clusters for transitioning the SnapMirror relationships.

[Preparing the cluster for transitioning volume SnapMirror relationships](#)

You must transition both the secondary and primary HA pairs in the same cutover window.

Steps

1. From the 7-Mode Transition Tool, perform a copy-free transition of the two HA pairs that contain the primary and secondary volumes of the 7-Mode volume SnapMirror relationship.

Before transitioning the 7-Mode HA pair that contains the secondary volumes, no manual intervention is required for the 7-Mode SnapMirror relationships. This ensures that the 7-Mode secondary volumes are transitioned as read-only volumes to ONTAP.

[Transitioning 7-Mode aggregates using copy-free transition](#)

2. Create an intercluster SVM peer relationship between the SVMs that contain the transitioned primary and

secondary volumes.

System administration

3. Create a volume SnapMirror relationship between the transitioned primary and secondary volumes.

Volume disaster recovery express preparation

4. On the destination volume, resynchronize the source volume and destination volume of the SnapMirror relationship:

snapmirror resync



At least one common Snapshot copy must exist between the source and destination volumes.

Troubleshooting: SnapMirror resynchronization fails if the common Snapshot copy was created in a Data ONTAP release earlier than 8.1. You can use the `-fs-version` parameter with the `volume snapshot show` command in the advanced privilege level to view the release in which the Snapshot copy was created. If you encounter this issue, break the SnapMirror relationship and then perform the resynchronization.

5. Monitor the status of the SnapMirror data transfers:

snapmirror show



You must not perform any operation, such as volume move or SnapMirror break, on the source and destination volumes until the resynchronization is completed successfully. Ensure that the resynchronization does not get aborted and completes successfully; otherwise, the volumes can go to an inconsistent state.

6. Commit the transition of the secondary project, followed by the primary project.

Committing the copy-free transition project

Related information

[ONTAP 9 commands](#)

[System Manager](#)

Troubleshooting transition issues

You should be aware of how to troubleshoot issues with the 7-Mode Transition Tool and where to look for log files. When you use the 7-Mode Transition Tool, you might see error messages that identify the issue and provide the solution.

Continuing with the transition if ignorable errors occur

During the transition, you might encounter some errors that block the transition. You can choose to ignore some of these errors by acknowledging the issues through the 7-Mode

Transition Tool CLI. You should rerun the failed operation after ignoring the error to continue with the transition.

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

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

Steps

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

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

ignorable_errorcategory is the type of error that you can ignore.

[Ignorable errors during transition](#)

- 2. Rerun the transition operation.

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

Ignorable errors during transition

You might encounter some ignorable errors during the transition. These errors can occur during the precheck, cabling, import, or commit operation of a copy-free transition project. You must acknowledge these errors before continuing with transition.

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

Precheck operation: ignorable error categories

Category	When the error is displayed
ignore-source-not-multipath	7-Mode disk shelves are not in a multipath configuration.
ignore-target-not-multipath	Disk shelves in the target cluster nodes are not in a multipath configuration.
ignore-source-storage-fault	7-Mode disk shelves have some fault (as displayed in the output of the <code>storage show fault</code> command).

Category	When the error is displayed
ignore-target-storage-fault	Disk shelves in the target cluster nodes have some fault (as displayed in the output of the <code>system node run -node node_name -command storage show fault</code> command).
ignore-target-port-requirement	Target cluster nodes do not have a sufficient number of ports available to connect the 7-Mode disk shelves.
ignore-aggr-space-less-than-5-percent	7-Mode aggregates are out of space because the free space in the 7-Mode aggregates is less than 5% of physical space.
ignore-aggr-logical-space-more-than-97-percent	7-Mode aggregates are out of space because the logical space in the aggregate is more than 97% full.
ignore-aggr-snapshot-spill-more-than-4-percent	7-Mode aggregates are out of space because Snapshot copies occupy more space than that allocated for the Snapshot copy reserve.
ignore-aggr-physical-space-more-than-89-percent-and-snapshot-spill	7-Mode aggregates are out of space because the total used physical space is more than 89% and the Snapshot copies occupy more space than that allocated for the Snapshot copy reserve.
ignore-volumes-with-file-guarantee	7-Mode volumes have space guarantee set to <code>file</code> , which is not supported in ONTAP.
ignore-volumes-with-disabled-guarantees	Space guarantee of volumes is currently disabled due to lack of space in the volumes.
nfs-qtrees-exported	<p>Qtree export rules are present in the 7-Mode system.</p> <p>Acknowledging this error means that you have understood the differences in the qtree export rules between Data ONTAP operating in 7-Mode and ONTAP. You might have to perform some manual steps after the NFS exports rules are applied by the 7-Mode Transition Tool.</p> <p>7MTT Precheck 10111 - Transitioning 7-Mode volumes that have qtree-level exports</p>
ignore-configuration-limits-check	<p>Objects and configurations to be transitioned exceed a certain limit. The storage cutover can take a long time and you must prepare for the downtime.</p> <p>Considerations for reducing the storage cutover time during Copy-Free Transition</p>

Category	When the error is displayed
ignore-cifs-ad-domain-mismatch	<p>The 7-Mode Transition Tool continues with the transition of the CIFS configuration even if the CIFS Active Directory domain of the 7-Mode system is different from the CIFS Active Directory domain of the target SVM.</p> <p>You must ensure that the CIFS Active Directory domains of the 7-Mode system and the target SVM are trusted domains. Otherwise, the transition of CIFS configurations to the target SVM fails.</p> <p>How to transition CIFS configurations when Active Directory Domain of CIFS server on 7-Mode and target SVM are different</p>

Cabling verification operation: ignorable error categories

Category	When the error is displayed
ignore-missing-spare-disks	One or more 7-Mode spare disks are not detected by the target cluster nodes.
ignore-missing-degraded-aggr-disks	<p>Cannot detect up to two disks from any of the 7-Mode RAID-DP RAID groups or one disk from any of the 7-Mode RAID-4 RAID groups on the target cluster nodes.</p> <p>Continuing with the transition makes such aggregates degraded after the aggregates are transitioned.</p>

Import operation: ignorable error categories

If you add an ignorable error category to the copy-free transition project during the import operation, Data ONTAP performs some corrective action on the aggregates and volumes in addition to changing the blocking error to a warning.

Category	When the error is displayed	Corrective action if the error is acknowledged and import operation is run again
ignore-aggregates-with-32bit-snapshot-for-import	32-bit Snapshot copies are detected in the 7-Mode aggregate.	32-bit Snapshot copies are deleted from all of the 7-Mode aggregates that are part of this project.
transition-dirty-aggregates-during-import	One of the transitioning aggregates was not cleanly shut down on the 7-Mode storage system.	All of the 7-Mode aggregates that were not cleanly shut down are transitioned. This might result in data loss after transition.

Category	When the error is displayed	Corrective action if the error is acknowledged and import operation is run again
ignore-aggregates-not-being-online-for-import	The aggregate was not online when the 7-Mode storage system was halted.	All of the offline aggregates are brought online.
ignore-volumes-with-32bit-snapshot-for-import	32-bit Snapshot copies are detected in the 7-Mode volume.	32-bit Snapshot copies are deleted from all of the 7-Mode volumes that are part of this project.
ignore-volumes-with-dirty-file-system-for-import	One of the transitioning volumes was not cleanly shut down on the 7-Mode storage system.	All of the 7-Mode volumes that were not cleanly shut down are transitioned. This might result in data loss after transition.
transition-offline-volumes-during-import	The volume was not online when the 7-Mode storage system was halted.	All of the offline volumes are brought online.
transition-restricted-volumes-during-import	The volume was in the restricted state when the 7-Mode storage system was halted.	All of the restricted volumes are brought online.

Commit operation: ignorable error categories

If you add an ignorable error category to the copy-free transition project during the commit operation, ONTAP performs some corrective action on the aggregates and volumes in addition to changing the blocking error to a warning.

Category	When the error is displayed	Corrective action if the error is acknowledged and commit operation is run again
ignore-commit-offline-aggregates	Some of the transitioned aggregates are offline.	All of the offline aggregates are brought online.

Downloading transition log files

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

Steps

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

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

Related information

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

Log files for the 7-Mode Transition Tool

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

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

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

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

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

Log file path	Contains information about...
transition-gui.log	Entries of all the actions performed by using the web interface
default/audit.log	<ul style="list-style-type: none">• All the parameters, such as HTTP or HTTPS port and log directory path, that are used by the tool every time 7-Mode Transition Tool is run• All the transition commands that are executed with the outputs
default/default/transition.log	Debug messages that are not specific to any project
default/STREAM_MANAGEMENT/stream_management.log	Debug messages that are logged by the scheduler while managing the schedules and which do not belong to any project
default/default/zapi-outbound.log	Output of all the Data ONTAP APIs that are executed by 7-Mode Transition Tool and which do not belong to any project
default/STREAM_MANAGEMENT/zapi-outbound.log	Output of all the Data ONTAP APIs that are executed by the 7-Mode Transition Tool scheduler while managing the schedules and which do not belong to any project

Log file path	Contains information about...
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.

Recovering from a failed LUN transition

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

Steps

1. Change to advanced privilege level:

```
set -privilege advanced
```

2. Check which LUNs failed:

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

3. Review the EMS logs and determine the corrective action that you must take.
4. Perform the required steps shown in the EMS message to correct the failure.
5. If any supported LUNs failed the transition, then to complete the transition:

```
lun transition start
```

6. View the transition status of the volumes:

```
lun transition show
```

The transition status can be one of following values:

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

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

Related information

[Space considerations when transitioning SAN volumes](#)

Failed to boot the 7-Mode controller in maintenance mode

The export and halt operation fails with the error message: Failed to boot the 7-Mode controller in maintenance mode. You must manually halt and boot the controller in the maintenance mode and rerun the operation.

Workaround

- 1. Halt the 7-Mode storage system:

```
halt -f -t 0
```

- 2. At the LOADER prompt, record the values set for the `bootarg.init.console_muted` and `bootarg.init.console_level` boot parameters:

```
printenv bootarg.init.console_muted

printenv bootarg.init.console_level
```

- 3. Disable the console messages by setting the following boot parameters:

```
setenv bootarg.init.console_muted "true"

setenv bootarg.init.console_level "-1"
```

- 4. From the 7-Mode Transition Tool, rerun the export and halt operation.
- 5. From the 7-Mode storage system, set the boot parameters to their original values, as recorded in Step 2:

If the boot parameters are...	Enter the following commands...
Not set previously with any value (undefined)	<pre>unsetenv bootarg.init.console_muted unsetenv bootarg.init.console_level</pre>
Set with a value previously	<pre>unsetenv bootarg.init.console_muted "original_value" unsetenv bootarg.init.console_level "original_value"</pre>

Performing a transition rollback to 7-Mode

Rollback refers to discontinuing the transition to ONTAP system and reverting to 7-Mode system. Transition rollback is manual; however, the 7-Mode Transition Tool provides the list of manual tasks you have to perform for a rollback.

You can roll back a transition project in the SVM provision, export, cabling, import, or preproduction testing phase. You cannot roll back after the transition project is committed.

Any new data written or modified, such as LUNs or LUN clones, in the transitioned volumes will be lost after a rollback. The volumes will be reverted to their original 7-Mode state.

When to roll back a transition and when to call technical support

You can roll back without assistance when performing a rollback on test or lab clusters, but you must call technical support if you encounter problems during or after transition, or if you want to roll back a transition performed on a production cluster.



You should not attempt to roll back a transition in a production environment without assistance from technical support.

If you encounter any of the following circumstances, contact technical support immediately:

- The transition process fails and cannot finish, and you are unsure of what to do next.
- The transition process finishes, but the cluster is unusable in a production environment.
- The transition process finishes and the cluster goes into production, but you are not satisfied with its behavior.
- The transition process finishes for some but not all of the data and configuration, and you decide that you want to roll back the transition.
- You have a problem with the transition process and cannot resolve the problem with the error response messages in the 7-Mode Transition Tool, in the Data ONTAP EMS event messages, and NetApp Knowledgebase.

[NetApp Knowledgebase](#)

Related information

[Downloading transition log files](#)

Rolling back a copy-free transition project

You can roll back a transition if you want to revert to 7-Mode at any stage of the copy-free transition before the transitioned aggregates are committed. Rollback is a manual operation. You can use the 7-Mode Transition Tool to generate the manual steps that you have to perform for a rollback.

- You must ensure that no volume or aggregate transition operations are running on the cluster.

You can use the `job show -jobtype transition` command.

- None of the 7-Mode aggregates must be committed.



You cannot perform the rollback if even one 7-Mode aggregate is committed.

- The target cluster nodes must not be in takeover mode.

Steps

1. Click **Rollback Prechecks** to verify that the project is eligible for a rollback.

If the precheck reports issues, you must fix them manually and rerun the precheck operation. For example, if you have created any new volumes or LUNs during preproduction testing, you must delete them manually.

2. Click **Generate Rollback Steps** to generate the list of manual steps that you have to perform for a successful rollback.
3. Click **Save As CSV** to save the manual steps in a file.

You can copy the rollback commands from the file and run them.

4. Depending on the transition phase from which you decide to roll back, perform the required manual steps:

- **Import or preproduction testing phase**

- i. Run the rollback commands on the cluster and click **Confirm**.
- ii. Connect the 7-Mode disk shelves to the 7-Mode controllers, verify the cabling manually, and click **Confirm**.
- iii. Run the rollback commands on the 7-Mode controllers and click **Confirm**.
- iv. View the configurations that are applied on the SVMs from the Operations History tab.
- v. Manually remove all configurations that were applied by the tool from the SVMs.

- **Cabling phase**

- i. Connect the 7-Mode disk shelves to the 7-Mode controllers, verify the cabling manually, and click **Confirm**.

You must ensure that the 7-Mode cabling matches the way it was at the start of the project.



You should use Config Advisor to verify the cabling.

- ii. Run the rollback commands on the 7-Mode controllers and click **Confirm**.
- iii. Manually remove all configurations that were applied by the tool from the SVMs.

You can view the configurations that are applied on the SVMs from the Operations History tab.

- **Export phase**

- i. Run the rollback commands on the 7-Mode controllers and click **Confirm**.
- ii. Manually remove all configurations that were applied by the tool from the SVMs.

You can view the configurations that are applied on the SVMs from the Operations History tab.

- **SVM provision phase**

Manually remove all configurations that were applied by the tool from the SVMs.

You can view the configurations that are applied on the SVMs from the Operations History tab.

[Manual steps for rolling back the transition](#)

5. After completing all the manual steps, click **Verify 7-Mode** from the 7-Mode Transition Tool to verify that the 7-Mode controllers are ready to serve data.

Manually rolling back the transition

You must perform some manual steps on the cluster and 7-Mode systems if you decide to roll back the transition. The list of manual rollback steps is generated by the 7-Mode Transition Tool.

The rollback steps vary depending on the stage at which you decide to roll back. You must perform all steps in this task if you decide to roll back after a successful import operation. If you decide to roll back at an earlier stage, you have to perform a subset of these steps.

Steps

1. Log in to the cluster.
2. If any transitioned volume is in a SnapMirror relationship, choose one of the following actions:
 - If a transitioned volume is the destination of a SnapMirror relationship, delete the SnapMirror relationship:
snapmirror delete -destination-path destination-path -source-path source-path
 - If a transitioned volume is the source of a SnapMirror relationship, release the SnapMirror relationship:
snapmirror release -destination-path destination-path -source-path source-path
3. From the cluster, verify that the following operations are not running on the transitioned volumes:
 - a. Volume move operation:
volume move show
 - b. LUN move operation:
lun move show
 - c. LUN copy operation:
lun copy show
4. Perform the rollback for all of the 7-Mode aggregates:
 - a. Log in to the diagnostic privilege level:
set -privilege diagnostic
 - b. Revert the aggregates to the 7-Mode state by using the `storage transition revert start` command.

This command requires additional parameters, such as the transition project ID and aggregate attributes. You should use the complete command with the parameters and its values, as generated by the 7-Mode Transition Tool.
 - c. Verify that the rollback is successful for all of the transitioned aggregates:
storage transition revert show-status

The `status-code` field for an aggregate is displayed as `revert_complete` when the rollback is successful.
5. Reassign the disk ownerships from the target cluster nodes to the 7-Mode controllers:

- a. Assign disk ownership to the 7-Mode controllers:

```
disk assign -disk disk_id -s system_id -force true
```

- b. Verify that the disk ownership is assigned to the 7-Mode controllers:

```
storage disk show -fields owner-id
```

6. Remove the 7-Mode LIFs from the SVMs:

```
network interface delete -vserver svm_name -lif lif_name
```

7. Remove the copy-free transition restrictions on the target cluster nodes from the diagnostic privilege level:

```
storage transition pre-commit end -session-id transition_project_id
```

You can also perform this step after the rollback operation is completed and the 7-Mode controllers are operational.

8. Delete the transition project information about the target cluster nodes by using the following diagnostic privilege level command:

```
storage transition purge-info -session-id transition_project_id
```

You can also perform this step after the rollback operation is completed and the 7-Mode controllers are operational.

9. If disk ownership autoassignment was disabled on the cluster nodes during the export and halt operation, enable it:

```
storage disk assign -auto true
```

10. Manually remove all of the configurations that the tool has transitioned to the target SVMs.

You can view the results of the SVM provision and import operations for information about the configurations transitioned by the tool.

11. Remove the disk shelves from the target cluster nodes, and then reconnect them to the 7-Mode controllers.



You should use the Config Advisor tool to verify the cabling.

12. If any 7-Mode disk shelf IDs were changed to resolve the conflicts with the disk shelf IDs of the target cluster nodes, manually change them to the old IDs and power-cycle the disk shelves for the new IDs to take effect.

13. Boot the source 7-Mode controllers in to the normal mode.

14. From one of the source 7-Mode controllers, enable the takeover capability:

```
cf enable
```

15. If automatic deletion of aggregate Snapshot copies was disabled during the export and halt operation, enable it:

```
options snap autodelete aggr_name on
```

Copyright information

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

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

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

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

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

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

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

Trademark information

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