



Transitioning 7-Mode aggregates using copy-free transition

ONTAP 7-Mode Transition

NetApp
October 09, 2025

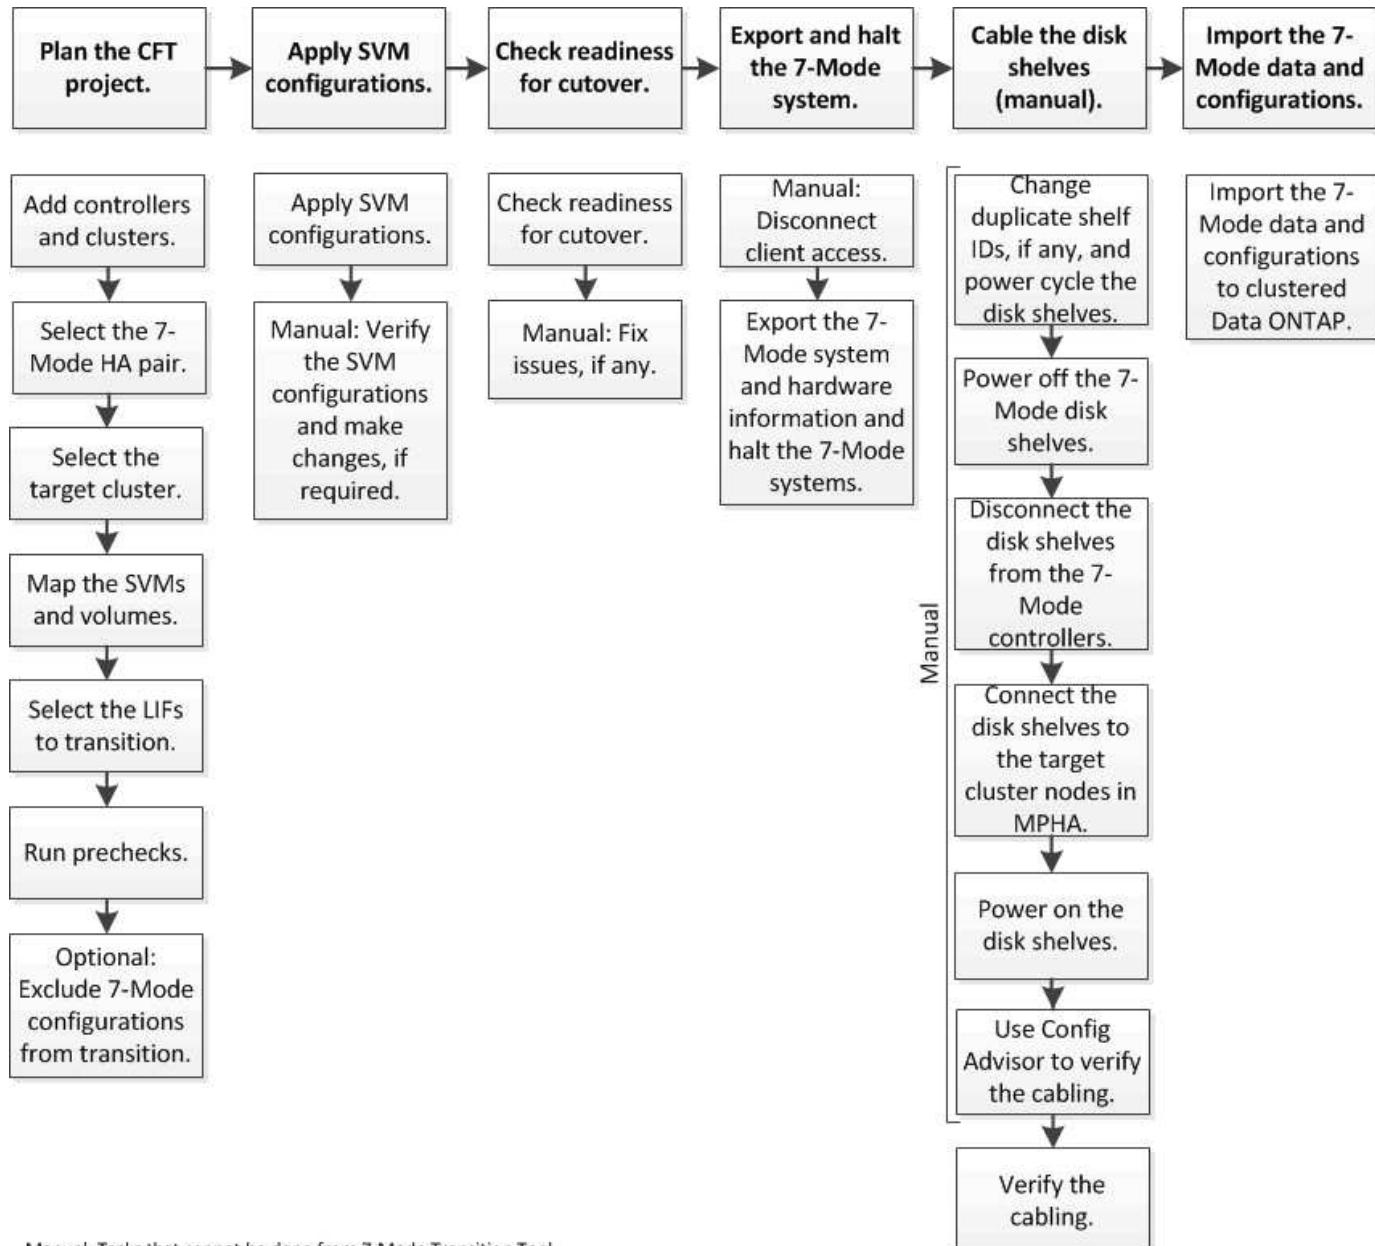
This PDF was generated from https://docs.netapp.com/us-en/ontap-7mode-transition/copy-free/reference_copy_free_transition_project_planning_worksheet.html on October 09, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Transitioning 7-Mode aggregates using copy-free transition	1
Planning a copy-free transition project	2
Copy-free transition project planning worksheets	2
Adding controllers and clusters	5
Creating a copy-free transition project	5
Selecting the target cluster nodes for transition	6
Mapping SVMs and volumes	7
Selecting LIFs for transition	8
Customizing the transition of 7-Mode configurations	9
Running prechecks	14
Applying SVM configurations	16
Verifying that 7-Mode systems are ready for cutover	17
Exporting storage configurations and halting 7-Mode systems	18
Disconnecting disk shelves from the 7-Mode system and connecting to cluster nodes	19
Importing 7-Mode data to ONTAP	21

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.



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

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

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

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

SVM	New IP address	Netmask	Default gateway	Home node	Home port
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

- 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.

- Click **Swap Node Mapping** to change the automatic assignment of source-to-target node mapping.
- 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	<ol style="list-style-type: none">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.

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

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

Example of consolidating Snapshot policies for transition

Snapshot schedules in 7-Mode and ONTAP before transition

7-Mode schedule

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

Snapshot policies existing in ONTAP

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

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

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

Snapshot policy in ONTAP after transition with consolidation and reuse

7-Mode volume	7-Mode Snapshot schedule	Snapshot policy in ONTAP
vol1	0 2 4@8,12,16,20 (weekly Snapshot copies: 0, daily Snapshot copies: 2, hourly Snapshot copies: 4 at 2, 4, 8, 12, 16, 20 hours)	Consolidated policy for vol1, vol2, and vol3 for which the existing ONTAP policy is reused Name: ScheduleDailyHourly4
vol2	0 2 4@8,12,16,20	vol3
0 2 4@8,12,16,20	vol4	1 2 3@8,12,16 (weekly Snapshot copies: 1, daily Snapshot copies: 2, hourly Snapshot copies: 3 at 8,12,16 hours)
<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	<p>Connect the disk shelves in a new stack to the target cluster nodes in a multipath configuration.</p> <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>
Ports are not available to connect the disk shelves	<p>Perform one of the following actions:</p> <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.

Copyright information

Copyright © 2025 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.