



7-Mode Transition Tool 1.2

Data and Configuration Transition Guide

For Transitioning to Clustered Data ONTAP®



Contents

Transitioning from 7-Mode to clustered Data ONTAP	7
What the 7-Mode Transition Tool does	7
Features and volumes not supported for transition	8
Types of transitions	9
How transition works	10
Preparing for transition	12
Data ONTAP version requirements for transition	12
License requirements for transition	12
Considerations for planning transition operations	13
Best practices for transition operations	14
Restrictions for transition	14
Preparing the cluster for transition	15
Preparing the 7-Mode system for transition	18
Port requirements for communicating with the 7-Mode Transition Tool	19
Prerequisites for transitioning NFS and CIFS configurations	20
Transition preparation checklist	22
Installing, upgrading, and uninstalling the 7-Mode Transition Tool	27
System requirements for 7-Mode Transition Tool	27
Installing the 7-Mode Transition Tool	28
Upgrading the 7-Mode Transition Tool	29
Uninstalling the 7-Mode Transition Tool	30
Understanding the 7-Mode Transition Tool	32
Transition terminology	32
Types of transition subprojects or sessions	33
How you transition a stand-alone volume	33
How you transition volumes in a SnapMirror relationship	36
Mode types you can use for completing the transition	41
States of a transition subproject or session	42
Considerations for using SnapMirror for transition	45
Guidelines for deciding when to perform cutover	47
Limits for transition	48
What the transition configuration file is	49

Severity levels for precheck messages	49
Considerations for transition	51
Considerations for NFS	51
Considerations for NFS exports transition	51
NFS only configuration transition	55
Considerations for CIFS	55
Considerations for transitioning CIFS local users and groups	55
Transitioning volumes using the 7-Mode Transition Tool (web interface)	57
Understanding the transition process	57
Project creation workflows	58
Steps to transition stand-alone volumes	58
How you create a project that transitions SnapMirror relationships	60
Understanding the project start and completion workflows	61
Dashboard	61
How you start and complete a project that transitions stand-alone volumes	63
How you start and complete a project that transitions volumes with SnapMirror relationships	64
Requirements and restrictions for using the web interface	66
Restrictions on creating and configuring projects	66
Prerequisites	68
Logging in to the tool	69
Creating and running a transition project	69
Adding controllers and clusters	70
Assessing controllers for transition	71
Creating a transition project	72
Starting a subproject	75
Completing a subproject	76
Managing a transition project	77
Running prechecks	77
Editing a subproject	78
Managing logical interfaces	79
Removing a volume or LIF from a subproject	80
Pausing and resuming a subproject	80
Aborting a subproject	81

Deleting a project	82
Collecting tool logs	82
Troubleshooting web interface issues	83
You cannot add credentials to controllers or clusters during transition	83
Cannot select a secondary volume from the Volume selection pane in some scenarios	83
Cannot select a volume with a broken SnapMirror relationship for transition	84
Cannot select a volume for transition if the tool fails to retrieve the volume information	84
Cannot proceed from the Volume mapping pane if the selected Vserver has no aggregate	84
Windows Program Compatibility Assistant notification displayed while installing the tool	84
Transitioning volumes using the 7-Mode Transition Tool (CLI)	85
Scenarios for selecting volumes in a session	85
Transitioning stand-alone volumes by using the transition wizard	87
Transition setup worksheet for a stand-alone volume	100
Transitioning a volume SnapMirror relationship using the transition wizard	103
Transition setup worksheet for a SnapMirror relationship	103
Transitioning secondary volumes by using the transition wizard	107
Transitioning primary volumes by using the transition wizard	119
Managing transition sessions by using the command-line interface	130
Creating a transition session by using the command-line interface	131
Running precheck on a session	135
Starting data copy for a transition session	136
Setting the audit log path of the Vserver	138
Pausing and resuming a copy operation	138
Completing the transition	139
Displaying transition session details	140
Aborting a transition session	141
Deleting a transition session	141
Viewing transition jobs	141
Commands for managing transition	143
Troubleshooting transition issues	146
Log files for the 7-Mode Transition Tool	146

Error messages in 7-Mode Transition Tool	147
Network connectivity issues with the 7-Mode system	152
Network connectivity issues with the cluster	154
Windows system out of memory	158
Session state changes to copy-unknown	159
Final update failed for one or more volumes during transition complete ...	161
Transition pause operation failed	164
Transition DR relationship not established if relationship is established using vFiler context	165
Installation or uninstallation of the 7-Mode Transition Tool is blocked	165
Error codes and corrective tasks for transition	167
Impact of Data ONTAP features on transition	176
Impact of takeover and giveback on transition	176
About backing up volumes in transition to tape	176
Glossary	178
Copyright information	185
Trademark information	186
How to send your comments	187
Index	188

Transitioning from 7-Mode to clustered Data ONTAP

The 7-Mode Transition Tool enables copy-based transitions of Data ONTAP 7G and 7-Mode FlexVol volumes and configurations to new hardware that is running clustered Data ONTAP version 8.2.0 or 8.2.1, with minimum client disruption and retention of storage efficiency options.

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

The 7-Mode Transition Tool enables you to:

- Migrate data and configuration from systems running Data ONTAP operating in 7-Mode to clustered Data ONTAP version 8.2.0 and clustered Data ONTAP 8.2.1.
- Transition existing FlexVol volumes from a 7-Mode controller to a Storage Virtual Machine (known as Vserver) on clustered Data ONTAP.
- Preserve Snapshot copies and storage efficiency in the volumes during the transition.
- Transition both the primary and the secondary FlexVol volumes in a SnapMirror relationship, and retain the SnapMirror relationship between the transitioned volumes that are on a cluster.
- Use the 7-Mode deployments while data is being copied to a newly deployed cluster.

Some 7-Mode configurations are not supported for transition, and some configurations behave differently after transitioning to clustered Data ONTAP. You must evaluate whether the existing 7-Mode configurations meet your requirements after they have been transitioned to a cluster. You must plan the transition process based on the following factors:

- 7-Mode features and functionality
- Workload (number of volumes that affect a single or group of related applications)
- Requirements after transitioning to clustered Data ONTAP

What the 7-Mode Transition Tool does

You can transition 7-Mode stand-alone volumes and volumes in a SnapMirror relationship by using the 7-Mode Transition Tool. The tool provides a web interface and automates the process of preparing for transition, copying data from multiple 7-Mode volumes, and applying the 7-Mode configuration to the Storage Virtual Machine (previously known as Vserver) .

The tool also allows you to assess if the 7-Mode controllers you have included in the tool for transition meet the key requirements for transition. This helps you screen the 7-Mode controllers (and all the vFiler units and volumes on them) for most of the common issues that prevent a successful transition.

You can also transition the volumes by using the 7-Mode Transition Tool CLI.

Related concepts

[*Understanding the 7-Mode Transition Tool*](#) on page 32

[*Transitioning volumes using the 7-Mode Transition Tool \(web interface\)*](#) on page 57

[*Transitioning volumes using the 7-Mode Transition Tool \(CLI\)*](#) on page 85

[*Features and volumes not supported for transition*](#) on page 8

Features and volumes not supported for transition

You cannot transition certain 7-Mode volumes, such as traditional volumes or SnapLock volumes, by using the 7-Mode Transition Tool. You also cannot transition certain 7-Mode features, such as synchronous SnapMirror relationships and MetroCluster configurations, because either some features might not be available in clustered Data ONTAP or the tool cannot transition these features.

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

Some of these features might require a manual post-transition setup in the cluster.

You cannot transition the following 7-Mode volumes:

- Volumes with LUNs and LUN clones (SAN transition)
- Restricted volumes
- Traditional volumes
- Volumes with NFS-to-CIFS character mapping (charmap)
- Volumes with Storage-Level Access Guard configurations
- SnapLock volumes
- FlexCache volumes
- 7-Mode volumes in a 64-bit aggregate to a 32-bit aggregate in clustered Data ONTAP
- FlexClone volumes can be transitioned; however, storage efficiency will be lost
- Root volume of a vFiler unit, where the root volume is based on a qtree that belongs to the default vFiler unit

You cannot transition the following 7-Mode configurations:

- IPv6 configurations
- SnapVault relationships
- FPolicy configurations
- Branch Cache configurations
- Antivirus configurations
- FCoE configurations
- PC-NFS
- WebNFS
- NFSv2

- Transition of sub-volume NFS exports (including qtree-level NFS exports) is not supported if the target cluster is running version 8.2.
Transition of qtree-level NFS exports is supported if the target cluster is running Data ONTAP 8.2.1. However, transition of sub-volume NFS exports other than qtree-level NFS exports is not supported if the target cluster is running Data ONTAP 8.2.1.
- Fencing of NFS export rules (fencing of NFS clients from a file system path)
- If the destination cluster is running clustered Data ONTAP 8.2, transition should not be attempted for 7-Mode volumes that are serving CIFS data and are being accessed by local users and groups because 7-Mode Transition Tool does not support the transition of local users and groups to clustered Data ONTAP 8.2.
- CIFS NT4 authentication
- CIFS NetBIOS aliases
- CIFS share-level ACLs with UNIX-style permissions
- Configuration to allow the administrator to connect to the CIFS home directories of other users (`cifs.homedirs_public_for_admin` option on the 7-Mode system)
- Volume SnapMirror cascade relationships
- Synchronous SnapMirror
- MetroCluster configurations
- Disaster recovery vFiler unit (vFiler DR)
Only the volumes in the source vFiler unit can be transitioned.
- NDMP configurations
- SFTP server
- TFTP server
- GARP VLAN Registration Protocol (GVRP)
- NIS for host name lookup
- Settings in the `/etc/hosts.equiv` file

Related concepts

[*Restrictions for transition*](#) on page 14

[*Considerations for transitioning CIFS local users and groups*](#) on page 55

Types of transitions

You can transition either stand-alone volumes or volumes that are in data protection relationships (in volume SnapMirror relationships).

You can transition the following:

- A group of 7-Mode stand-alone volumes
- A group of 7-Mode volumes in volume SnapMirror relationships, with their volume SnapMirror relationships.

You can transition multiple groups of volumes together. However, you need to consider and plan for the replication streams required by the 7-Mode Transition Tool in addition to the replication streams required by your existing backup and disaster recovery (DR) operations. For more information about how to manage the replication stream requirements of the 7-Mode Transition Tool, see [*Considerations for using SnapMirror for transition*](#) on page 45.

Related concepts

[*Types of transition subprojects or sessions*](#) on page 33

[*Mode types you can use for completing the transition*](#) on page 41

How transition works

The transition workflow consists of the following phases: the preparation phase, the data copy phase, the precutover phase, and the cutover phase. Understanding what occurs during each phase helps you to manage your transition efficiently.

1. Preparation phase

Information about the 7-Mode system and cluster is provided and a verification whether the configuration of 7-Mode volumes can be applied to the clustered Data ONTAP volumes is performed.

2. Data copy phase

The Vserver (if a new Vserver is selected) and its volumes are created and data copy starts from the 7-Mode volumes to the clustered Data ONTAP volumes.

A transition peer relationship is established between the 7-Mode system and the cluster. Then, a SnapMirror relationship is established between the 7-Mode volume and clustered Data ONTAP volume, and baseline transfer is started. After the baseline transfer is completed, updates occur according to the selected SnapMirror schedule.

3. Precutover phase

This phase requires manual intervention where you monitor the status of the data copy and prepare for cutover. You must plan a downtime for all clients that will be impacted by this transition.

4. Cutover phase

In this phase, configuration of the 7-Mode volumes is applied to the clustered Data ONTAP volumes. Also, you must disconnect client access from the 7-Mode volumes.

A final data transfer is performed from the 7-Mode volumes to the clustered Data ONTAP volumes and the SnapMirror relationship between the 7-Mode and clustered Data ONTAP volumes is broken. Data LIFs are configured on the Vserver. The configuration of the 7-Mode volumes and vFiler unit (if any) is collected, translated, and applied to the clustered Data ONTAP volumes and Vserver.

You can redirect client access to the clustered Data ONTAP volumes after the cutover phase is completed successfully.

Related concepts

Mode types you can use for completing the transition on page 41

States of a transition subproject or session on page 42

Types of transition subprojects or sessions on page 33

Related references

Limits for transition on page 48

Preparing for transition

Before initiating a data copy operation from 7-Mode to clustered Data ONTAP, you must understand the requirements and restrictions for transition, and complete certain tasks on the 7-Mode system and the cluster.

You must ensure that the following requirements are met before you transition the 7-Mode volumes:

- The 7-Mode systems must be running Data ONTAP versions that are supported for transition.
- SnapMirror must be licensed on the 7-Mode system for transition.
- Required feature licenses, if exists on the 7-Mode system, must be installed on the cluster.
- NTP server must be configured and time must be synchronized across the 7-Mode system and cluster.
- All preparatory tasks on the 7-Mode system must be completed.
- All preparatory tasks on the cluster must be completed.

Related concepts

[Transitioning from 7-Mode to clustered Data ONTAP](#) on page 7

[Understanding the 7-Mode Transition Tool](#) on page 32

Data ONTAP version requirements for transition

You should be aware of the versions of Data ONTAP operating in 7-Mode that are supported for transitioning to clustered Data ONTAP version 8.2.0 or 8.2.1.

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

Related concepts

[License requirements for transition](#) on page 12

Related information

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

License requirements for transition

Before you transition a volume from 7-Mode to clustered Data ONTAP, you must ensure that SnapMirror is licensed on the 7-Mode storage system. Additionally, SnapMirror licenses are required

on the primary and secondary clusters to transition a 7-Mode volume SnapMirror relationship. The 7-Mode Transition Tool does not require any license.

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

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

Related concepts

[Data ONTAP version requirements for transition](#) on page 12

Related information

[Documentation on the NetApp Support Site: *support.netapp.com*](#)

Considerations for planning transition operations

When you select the 7-Mode volumes and configurations for transition, you must consider how the transition might affect the clients and applications. You can limit the scope of a single transition operation to reduce the effect and downtime for applications.

1. When transitioning root volumes, you must transition root volumes only as part of the final subproject or session that contains the last set of volumes in a particular controller or a vFiler unit.
2. When transitioning FlexClone volumes to the Vserver, the clone hierarchy and storage efficiency are lost in the process.
3. For deduplication and compression to take effect on any new data written on the clustered Data ONTAP volume, you must enable deduplication and compression schedules after the transition.

Note: If you have enabled deduplication and compression on the 7-Mode volumes, the existing data in the transitioned volumes is deduplicated and compressed.

4. If multiple applications and clients access the same resources (for example, a volume, group of volumes, or a vFiler unit) on the storage controller, you have to appropriately select the volumes to transition.

For example, if you are using vFiler units (MultiStore) for workload isolation and resource sharing among clients and applications, you can transition the volumes (up to a maximum of 20 volumes) in the vFiler unit in a single transition operation by using the 7-Mode Transition Tool. The following are other situations you might consider when planning a transition operation:

- Multiple workloads that access the same volume or set of volumes
- Workloads that share the same NFS exports or CIFS shares
- Several workloads that share the same IP address to access the storage system

Related concepts

[Data ONTAP version requirements for transition](#) on page 12

[License requirements for transition](#) on page 12

[Features and volumes not supported for transition](#) on page 8

[Restrictions for transition](#) on page 14

Best practices for transition operations

You can follow the best practices and recommendations to manage the transition operations effectively and efficiently.

- You should not modify the objects (volumes, IP addresses, system information, and so on) after the subproject or session is prepared.
- You should avoid using multiple web interface sessions simultaneously to prevent undesired results.
- You should avoid modifying the controller and cluster passwords during the transition process.

Restrictions for transition

You must be aware of certain restrictions for transitioning some 7-Mode volumes and configurations such as FlexClone volumes and volumes with deduplication and compression.

- If you want to transition 7-Mode primary and secondary volumes when both the 7-Mode source and destination are running Data ONTAP 7.3.x or 8.0.x, you must start transitioning the 7-Mode secondary volume only when there are no data updates from the 7-Mode primary to the 7-Mode secondary volume.

You must ensure that the data update schedules for the 7-Mode primary volume to the 7-Mode secondary volume do not conflict with the schedules for the 7-Mode secondary volume to the clustered Data ONTAP secondary volume.

- The tool does not transition a vFiler unit as a single entity.
However, you can transition all the volumes in a vFiler unit by selecting them as a part of one or more subprojects or sessions.
- The tool does not transition root volume of a vFiler unit if the root volume is based on a qtree that belongs to the default vFiler unit.
- The tool does not allow you select a Vserver that does not have the Network File System (NFS) or the Common Internet File System (CIFS) protocol enabled.
- You must not initiate transition while the aggregates on either the 7-Mode system or cluster are upgrading from 32-bit to 64-bit format; otherwise transition fails.
- The tool does not transition a volume with a qtree if the volume and qtree are owned by different vFiler units.

Transitioning such a volume will cause the qtree to become inaccessible.

- The tool does not transition a volume with a qtree that is the destination of a qtree SnapMirror relationship.
The qtree SnapMirror relationship must be broken before the volume can be transitioned.
- 7-Mode Snapshot copies that have LUNs cannot be used to restore a clustered Data ONTAP volume.
- You cannot transition a fanout SnapMirror relationship (a primary volume that is in SnapMirror relationships with more than one secondary volume in different controllers) by using the 7-Mode Transition Tool (web interface).
To transition the SnapMirror relationships in a fanout configuration, you must use the 7-Mode Transition Tool (CLI). You should create separate sessions for each secondary volume, complete the transition of the secondary sessions, and then create and complete the transition of the primary volume.
- You cannot transition volumes from different vFiler units or from different 7-Mode controllers to the same Vserver at the same time. You must complete the transition of volumes from a given vFiler unit or 7-Mode controller before you can start the transition of volumes from another vFiler unit or 7-Mode controller.

The precheck operation displays information about some of these restrictions.

Related concepts

[*Features and volumes not supported for transition*](#) on page 8

Preparing the cluster for transition

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

Before you begin

- The cluster must already be set up.
For information about setting up a cluster, see the *Clustered Data ONTAP Software Setup Guide*.
- The cluster must be healthy and none of the nodes must be in takeover mode.
- The aggregates that are used to place the transitioned volumes in the Vserver must have an SFO policy.
- The aggregates must be on nodes that have not reached the maximum volume limit.
- If you want to transition volumes from a 32-bit aggregate of a 7-Mode system to a 64-bit aggregate of a cluster, you must have provided an additional 5 percent space in the destination aggregate.
The additional space is required to upgrade the transitioned volume to 64-bit format. For more information about aggregates, see the *Clustered Data ONTAP Physical Storage Management Guide*.

- For establishing a Vserver peer relationship when transitioning a volume SnapMirror relationship, the following conditions must be met:
 - The secondary cluster should not have a Vserver with the same name as that of the primary Vserver.
 - The primary cluster should not have a Vserver with the same name as that of the secondary Vserver.

Steps

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


```
ssh username@cluster_mgmt_IP
```
2. Use the `system services web show` command to verify that SSL is enabled on the cluster.

Example

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

3. Verify that HTTPS is allowed on the cluster-management LIF by performing the following steps:
 - a. Use the `network interface show` command to view the firewall policy for the cluster-management LIF.

Example

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

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

Example

```
cluster1::> system services firewall policy show -policy mgmt
Policy      Service      Action IP-List
-----
mgmt
```

```

      dns      allow  0.0.0.0/0, ::/0
      http     allow  0.0.0.0/0, ::/0
      https    allow  0.0.0.0/0, ::/0
      ndmp     allow  0.0.0.0/0, ::/0
      ntp      allow  0.0.0.0/0, ::/0
      rsh      deny   0.0.0.0/0, ::/0
      snmp     allow  0.0.0.0/0, ::/0
      ssh      allow  0.0.0.0/0, ::/0
      telnet   deny   0.0.0.0/0, ::/0
9 entries were displayed.

```

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

4. Create the intercluster LIF on each node of the cluster for communication between the cluster and 7-Mode system by performing the following steps:

- a. Use the `network interface create` command to create an intercluster LIF.

Example

```

cluster1::> network interface create -vserver cluster1-01 -lif
intercluster_lif -role intercluster -home-node cluster1-01 -home-
port e0c -address
192.0.2.130 -netmask 255.255.255.0

```

- b. Use the `network routing-groups route create` command to create a static route for the intercluster LIF.

Example

```

cluster1::> network routing-groups route create -vserver
cluster1-01 -routing-group i192.0.0.0/18 -destination 0.0.0.0/0 -
gateway 192.0.2.129

```

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

Example

```

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

```

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

Related tasks

[Preparing the 7-Mode system for transition](#) on page 18

Related information

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

Preparing the 7-Mode system for transition

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

Before you begin

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

About this task

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

Steps

1. Add and enable the SnapMirror license on the 7-Mode system.
 - a. Add the SnapMirror license on the 7-Mode system by using the following command:
license add xxxxxxxx
xxxxxxx is the license code you purchased.
 - b. Enable SnapMirror by entering the following command:
options snapmirror.enable on
 - c. Configure the 7-Mode system and the target cluster to communicate with each other by performing one of the following actions:
 - Set the `snapmirror.access` option to `all`.
 - Set the value of the `snapmirror.access` option to the IP addresses of all the intercluster LIFs on the cluster.
 - If the `snapmirror.access` option is `legacy` and the `snapmirror.checkip.enable` is off, add the Vserver name to the `/etc/snapmirror.allow` file.
 - If the `snapmirror.access` option is `legacy` and the `snapmirror.checkip.enable` is on, add the IP addresses of the intercluster LIFs to the `/etc/snapmirror.allow` file.
2. Enable SSL on the 7-Mode system by performing the following steps:
 - a. Set up and start SSL by entering the following command:
secureadmin setup ssl
 - b. Enable SSL by entering the following command:
options ssl.enable on
 - c. If you are running Data ONTAP 8.x, enable SSLv3 by entering the following command:

```
options ssl.v3.enable on
```

3. Depending on the Data ONTAP version of your 7-Mode system, perform the following steps:

a. Allow SnapMirror traffic on all the interfaces by entering the following command:

```
options interface.snapmirror.blocked ""
```

b. If you are running Data ONTAP versions 7.3.7, 8.0.3, or 8.1 and using the IP address of the e0M interface as the management IP address to interact with the 7-Mode Transition Tool, allow data traffic on the e0M interface by entering the following command:

```
options interface.blocked.mgmt_data_traffic off
```

4. If you have set these volume-level options on any volumes to be transitioned, disable the options after ensuring that disabling them does not impact other operations:

```
vol options vol_name no_i2p off
```

```
vol options vol_name read_realloc off
```

```
vol options vol_name nvfail off
```

Related tasks

[Preparing the cluster for transition](#) on page 15

Related information

[Documentation on the NetApp Support Site: support.netapp.com](#)

Port requirements for communicating with the 7-Mode Transition Tool

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

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

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

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

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

Related concepts

[System requirements for 7-Mode Transition Tool](#) on page 27

Prerequisites for transitioning NFS and CIFS configurations

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

Transitioning NFS

- On the 7-Mode system, the following conditions must be met:
 - NFS must be licensed.
 - If MultiStore license is enabled, NFS must be added to the list of allowed protocols for the vFiler unit that owns the transitioning volumes.
 - For transitioning an Active Directory (AD) based Kerberos server to a new Vserver, a DNS entry must exist for the AD domain.
- On the cluster, the following conditions must be met:
 - NFS must be licensed.
 - NFS must be added to the list of allowed protocols for the Vserver.
 - For transitioning an AD-based Kerberos server to an existing Vserver with DNS configured, a DNS entry must exist for the AD domain.
 - The clock skew between the Kerberos key distribution center (KDC) and the clustered Data ONTAP system must be less than or equal to 5 minutes.

Transitioning CIFS

- On the 7-Mode system, the following conditions must be met:
 - CIFS license must be added.
 - If MultiStore license is enabled, CIFS must be added to the list of allowed protocols for the vFiler unit that owns the transitioning volumes.
 - CIFS must be set up and running.
 - The authentication type for CIFS must be Active Directory (AD).
 - At least one volume that is selected for transition must be configured for serving CIFS requests.
- On the cluster, the following conditions must be met:
 - CIFS license must be added.
 - CIFS must be added to the list of allowed protocols for the Vserver.
 - DNS must be configured for the Vserver.
 - CIFS server must be manually configured before cutover.
 - CIFS must be running.

After disconnecting client access to the volumes, you can create the CIFS server on the Vserver before cutover in the following two ways:

- Transfer or preserve identity: You can create the CIFS server on the Vserver by reusing the same CIFS server name of the 7-Mode system or vFiler unit.
To preserve the identity of the CIFS server, the following conditions must be met:
 - The number of volumes in the 7-Mode system or vFiler unit to be transitioned must be less than or equal to 20 so that all the 7-Mode volumes can be transitioned in a single subproject or session.
 - Before cutover, the CIFS server must be re-created on the 7-Mode system with a different name.
 - During cutover, there must not be any active CIFS sessions on the 7-Mode system.
 - During cutover, the new CIFS server must be up and running on the 7-Mode system.
- Use a new identity: You can create the CIFS server on the Vserver by using a new CIFS server name.
The CIFS server must be up and running on the 7-Mode system.

Related concepts

Considerations for NFS exports transition on page 51

Considerations for transitioning CIFS local users and groups on page 55

Transition preparation checklist

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

- [Data ONTAP version requirements](#) on page 22
- [Licensing requirements](#) on page 23
- [SnapMirror requirements on the 7-Mode system](#) on page 23
- [Volume options on the 7-Mode system](#) on page 23
- [Managing access to the cluster](#) on page 24
- [Managing access to the 7-Mode system](#) on page 24
- [Networking requirements](#) on page 24
- [Port requirements](#) on page 25

Data ONTAP version requirements

Item	Yes
7-Mode version is one of the following: • Data ONTAP 7.3.3 • Data ONTAP 7.3.4 • Data ONTAP 7.3.5 • Data ONTAP 7.3.6 • Data ONTAP 7.3.7 • Data ONTAP 8.0.2 • Data ONTAP 8.0.3 • Data ONTAP 8.0.4 • Data ONTAP 8.0.5 • Data ONTAP 8.1.0 • Data ONTAP 8.1.2 • Data ONTAP 8.1.3	
Your cluster must be running one of the following Data ONTAP versions: • Data ONTAP 8.2 • Data ONTAP 8.2.1	

For the updated list of Data ONTAP versions supported for transition by the 7-Mode Transition Tool, see the Interoperability Matrix

Licensing requirements

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

SnapMirror requirements on the 7-Mode system

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

Volume settings on the 7-Mode system

Item	Yes
Volume is online	
Volume is not restricted	

Item	Yes
The following volume options are disabled: • no_i2p • read_realloc • nvfail	

Managing access to the cluster

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

Managing access to the 7-Mode system

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

Networking requirements

Item	Yes
Cluster is reachable using the cluster-management LIF	
One or more intercluster LIFs are set up on each node of the cluster For multipathing, two intercluster LIFs are required on each node.	
Static routes are created for the intercluster LIFs	
7-Mode system and cluster are reachable from the Windows system on which the 7-Mode Transition Tool is installed	

Item	Yes
NTP server is configured and the 7-Mode system time is synchronized with the cluster time	

Port requirements

Item	Yes
7-Mode system • 10565/TCP • 10566/TCP • 10567/TCP • 10568/TCP • 10569/TCP • 10670/TCP • 80/TCP • 443/TCP	
Cluster • 10565/TCP • 10566/TCP • 10567/TCP • 10568/TCP • 10569/TCP • 10670/TCP • 11105/TCP • 80/TCP • 443/TCP	

Related concepts

[Restrictions for transition](#) on page 14

[Data ONTAP version requirements for transition](#) on page 12

[License requirements for transition](#) on page 12

Related tasks

[Preparing the 7-Mode system for transition](#) on page 18

[Preparing the cluster for transition](#) on page 15

Related information

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

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

You can download and install the 7-Mode Transition Tool on a system that is running a 64-bit Windows 7 operating system. You can also upgrade the tool from the 1.1 version to the 1.2 version. When required, you can uninstall the 7-Mode Transition Tool.

Related concepts

[Transitioning from 7-Mode to clustered Data ONTAP](#) on page 7

[Preparing for transition](#) on page 12

[Understanding the 7-Mode Transition Tool](#) on page 32

System requirements for 7-Mode Transition Tool

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

Hardware requirements

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

Software requirements

- 64-bit Windows 7, Windows 7 Service Pack 1 (SP1), or Windows Server 2008 R2 SP1
- 64-bit Oracle Java Runtime Environment (JRE) 1.6.0.32

Note: If the Windows host does not have JRE or has 32-bit JRE installed, the installer for the 7-Mode Transition Tool automatically installs 64-bit JRE 1.6.0.32. If the host has an earlier version of 64-bit JRE, the installer automatically updates it to JRE 1.6.0.32.

Server (web interface) configuration requirements

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

- Ports 8443 (for HTTPS) and 8088 (for HTTP) must be available.
If port 8443 is not available, or if you want to use a different port for HTTPS, you must change the port specified by the `tool.https.port` parameter in the `transition-tool.conf` file.

For an alternative HTTP port, you must change the port specified by the `tool.http.port` parameter in the `transition-tool.conf` file.

Note: You must restart the NetApp Transition Tool service after changing the port in the configuration file.

- The firewall should be either turned off or configured to allow traffic on the port that is used to access the tool.

Client (web interface) requirements

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

- One of the following web browsers:
 - Microsoft Internet Explorer 9.0 or 10
 - Google Chrome 27 or later
 - Mozilla Firefox 20 or later
- A screen resolution of 1280 × 1024 or higher

For more information, see the Interoperability Matrix.

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

Related concepts

[Port requirements for communicating with the 7-Mode Transition Tool](#) on page 19

Related tasks

[Installing the 7-Mode Transition Tool](#) on page 28

[Uninstalling the 7-Mode Transition Tool](#) on page 30

Related information

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

Installing the 7-Mode Transition Tool

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

Before you begin

- You must have reviewed the *[System requirements for the 7-Mode Transition Tool](#)*.
- You must have the necessary administrator privileges to install and launch the application.
- You must have downloaded the .exe file from the NetApp Support Site.

- Every time you install a new version of the tool, it is necessary to clear the browser cache by pressing Ctrl+F5.
- If you have 1.0 version of the 7-Mode Transition Tool, you must first uninstall it before you install 7-Mode Transition Tool 1.2.

Attention: You must complete any in-progress sessions or subprojects before you uninstall the earlier version of the tool.

- If you uninstall an earlier version but choose to retain the log files, you must take a backup of the log files that you require before you start the 7-Mode Transition Tool 1.2 installation. The 7-Mode Transition Tool 1.2 installer installs the tool to the existing folder; any files still found in that folder will be overwritten.

Steps

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

Related concepts

[Data ONTAP version requirements for transition](#) on page 12

[License requirements for transition](#) on page 12

[Port requirements for communicating with the 7-Mode Transition Tool](#) on page 19

[Understanding the 7-Mode Transition Tool](#) on page 32

Related information

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

Upgrading the 7-Mode Transition Tool

You can upgrade the 7-Mode Transition Tool on your Windows system. The 7-Mode Transition Tool upgrade is supported only from version 1.1 to version 1.2.

Before you begin

- You must have reviewed the [System requirements for the 7-Mode Transition Tool](#).
- You must have the necessary administrator privileges to install and launch the application.
- The projects must be in completed or in an aborted state in version 1.1.
- You must make sure that all instances of the tool are closed before beginning the upgrade.
- You cannot upgrade the 7-Mode Transition Tool from the 1.0 version to the 1.1 version or the 1.2 version.

Steps

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

Related information

NetApp Support Site: support.netapp.com

Uninstalling the 7-Mode Transition Tool

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

Before you begin

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

About this task

JRE is not removed as part of the uninstallation.

Steps

1. Click **Start > Control Panel > Uninstall a program**.
2. Select the 7-Mode Transition Tool from the list of programs and click **Uninstall**.
3. Complete the wizard to uninstall the 7-Mode Transition Tool.

You can choose to save the session details and logs for troubleshooting purposes.

After you finish

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

Related tasks

Installing the 7-Mode Transition Tool on page 28

Understanding the 7-Mode Transition Tool

The 7-Mode Transition Tool enables you to transition 7-Mode volumes and configuration to clustered Data ONTAP with minimum client disruption. The 7-Mode Transition Tool provides both web-based graphical and command-line interfaces for performing the transition.

The tool interacts with the 7-Mode storage system to collect information about the volumes and IP address, verifies whether the 7-Mode configuration can be applied to clustered Data ONTAP system, and applies the configuration to the Vserver on the cluster.

You must use only one interface for an end-to end transition workflow. For example, if you start a transition operation for a set of volumes by using the web interface, you must complete the transition for those volumes by using the same interface.

Related concepts

[Transitioning volumes using the 7-Mode Transition Tool \(web interface\)](#) on page 57

[Transitioning volumes using the 7-Mode Transition Tool \(CLI\)](#) on page 85

Transition terminology

Project	In the 7-Mode Transition Tool (web interface), a project enables you to configure and manage the transition of a group of volumes.
Subproject	In the 7-Mode Transition Tool (web interface), a subproject is a container within a project that groups a set of volumes for transition. A subproject also contains the configuration details that are required for the transition, such as the Vserver mapping, volume mapping, interface mapping, and SnapMirror schedule.
Session	A transition session defines the source and destination of transition. It is the basic unit of transition that is used by the 7-Mode Transition Tool (CLI) to transition 7-Mode volumes. A transition session consists of 7-Mode objects (such as the host name, IP addresses, and volumes) and clustered Data ONTAP objects (such as a cluster, a Vserver, volumes, and LIFs). Note: The two terms—session and subproject—can be used interchangeably.
Transition peer relationship	A transition peer relationship is an authorization mechanism that enables SnapMirror to establish a relationship between a 7-Mode storage system and a Vserver, allowing data to be copied from the 7-Mode volumes to the clustered Data ONTAP volumes.

Types of transition subprojects or sessions

Transition subprojects or sessions are of three types: stand-alone, primary, and secondary. You should understand what these subproject or session types are to retain your data protection and configurations after transition.

Stand-alone subproject or session	Volumes in this subproject or session are not in any SnapMirror relationship with volumes in other storage systems.
Primary subproject or session	Volumes in this subproject or session are the source volumes of a SnapMirror relationship.
Secondary subproject or session	Volumes in this subproject or session are the destination volumes of a SnapMirror relationship.

You can execute only one of the following transition project at any given time:

- A stand-alone subproject or session
- One secondary subproject or session and one primary subproject or session

Note: To transition a 7-Mode SnapMirror relationship while retaining the data protection at all times, you must first prepare and complete the transition of volumes in the secondary session, followed by the primary session.

Related concepts


[*Mode types you can use for completing the transition*](#) on page 41

[*States of a transition subproject or session*](#) on page 42

How you transition a stand-alone volume

Transitioning a stand-alone volume includes different phases: preparation, data copy, precutover, and cutover. Understanding what occurs during each phase helps you manage your transition efficiently.

The following image illustrates the steps involved in the transition of a stand-alone volume:



Preparation phase

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

- Collecting and adding 7-Mode storage system and volume information
- Running the transition precheck
- Collecting and adding cluster and Vserver information
- Collecting IP addresses that must be configured on the Vserver:
 - Selecting IP addresses that exist on the 7-Mode system
 - Specifying new IP addresses that must be configured on the Vserver
- Collecting the SnapMirror schedule from the cluster for update transfers

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

Data copy phase

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

1. Disabling administrative access to the clustered Data ONTAP volumes, making them read-only
2. Setting up a transition peer relationship between the 7-Mode system and the Vserver
3. Establishing a SnapMirror relationship between the 7-Mode volumes and clustered Data ONTAP volumes
4. Completing the baseline transfer
5. Performing scheduled SnapMirror updates to the clustered Data ONTAP volumes

Note: You cannot edit a session or a subproject after starting the data copy operations.

Precutover phase

This is the preparation phase for the cutover. You must perform the following tasks:

1. Monitor the data copy progress and ensure that data copy is completed.
2. If you have CIFS configured on the 7-Mode system, manually set up the CIFS server on the Vserver before cutover.
3. You must plan downtime for all the clients that will be impacted by this transition.
4. Ensure that you have reviewed the information given in [Guidelines for deciding when to cutover](#) on page 47.

Cutover phase

In this phase, you must disconnect client access from the 7-Mode system. When you initiate the cutover, the following tasks are performed by the 7-Mode Transition Tool in this phase:

1. Performing a final SnapMirror update
2. Breaking and deleting the SnapMirror relationship
3. Collecting configuration from 7-Mode volumes and applying it to the volumes in the Vserver
4. Taking the 7-Mode volumes offline and removing the 7-Mode IP addresses selected for transition
5. Configuring the data LIFs on the Vserver
6. Enabling administrative access to the clustered Data ONTAP volume

Resuming client access

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

1. Create the DNS entry for the domain names, which will be generated from the domain name specified in

options dns.domainname

option on the 7-Mode system, using the following command:

```
dns create -vserver <vserver-name> -domains <comma-separated-domain-names-list> -name-server <comma-separated-name-servers-list>
```

Example: If

options dns.domainname

on 7-Mode is set to `eng.btc.netapp.in`, then the DNS entry should be configured on the cluster after the transition is complete using the following list of search domains.

```
dns create -vserver <vserver-name> -domains  
eng.btc.netapp.in,btc.netapp.in,netapp.in -name-server <name-servers>
```

2. Configure the features, such as IPv6 and FPolicy, that were available on the 7-Mode system, but were not transitioned automatically to the Vserver by the tool.
3. Ensure that the Vserver is ready to serve data to the clients by verifying the following:
 - The volumes on the Vserver are online and read/write.
 - The transitioned IP addresses are up and reachable on the Vserver.
4. Redirect client access to the clustered Data ONTAP volumes.

Related concepts

[*How you transition volumes in a SnapMirror relationship*](#) on page 36

[*Features and volumes not supported for transition*](#) on page 8

[*Restrictions for transition*](#) on page 14

How you transition volumes in a SnapMirror relationship

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

The SnapMirror schedule used by the 7-Mode volumes is applied to the SnapMirror relationships between the 7-Mode primary volumes and clustered Data ONTAP secondary volumes.

Preparation phase

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

1. Collecting and adding 7-Mode storage system and volume information
2. Collecting the following information for transitioning a volume SnapMirror relationship:
 - For transitioning a secondary volume, information about the 7-Mode primary system
 - For transitioning a primary volume, information about the secondary Vserver
3. Running the transition precheck
4. Collecting and adding cluster and Vserver information
5. Collecting IP addresses that must be configured on the Vserver:
 - Selecting IP addresses that exist on the 7-Mode system
 - Specifying new IP addresses that must be configured on the Vserver
6. Collecting the SnapMirror schedule from the clusters for update transfers

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

Data copy phase

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

1. Disabling administrative access to the clustered Data ONTAP volume
2. Setting up a transition peer relationship between the 7-Mode system and the Vserver
3. Establishing a SnapMirror relationship between the 7-Mode volumes and clustered Data ONTAP volumes
4. Completing the baseline transfer
5. Performing scheduled SnapMirror updates to the clustered Data ONTAP volumes

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

Precutover phase

This is the preparation phase for the cutover. You must perform the following tasks:


1. Monitor the data copy progress and ensure that data copy is completed.
2. If you have CIFS configured on the 7-Mode system, manually set up the CIFS server on the Vserver before cutover.
3. You must plan downtime for all the clients that will be impacted by this transition.
4. Ensure that you have reviewed the information given in [Guidelines for deciding when to cutover](#) on page 47.

Cutover phase (secondary)

In this phase, you must disconnect client access from the 7-Mode secondary system. When you initiate cutover for the secondary volume, the 7-Mode Transition Tool performs the following tasks in this phase:

1. Setting up a transition peer relationship with the 7-Mode primary system
2. Performing a final SnapMirror update from the 7-Mode secondary volume to the clustered Data ONTAP secondary volume
3. Breaking and deleting the SnapMirror relationship between the 7-Mode secondary volume to the clustered Data ONTAP secondary volume
4. Enabling administrative access to the clustered Data ONTAP secondary volume
5. Applying volume and services configuration to the Vserver
6. Establishing a SnapMirror relationship between the volumes on the 7-Mode primary system to the clustered Data ONTAP secondary volumes
7. Taking the 7-Mode volumes offline and removing the 7-Mode IP addresses selected for transition
8. Configuring the data LIFs on the Vserver

The following illustration depicts the transition of a secondary volume:




Cutover phase (primary)

In this phase, you must disconnect client access from the 7-Mode primary system. When you initiate cutover for the primary volume, the 7-Mode Transition Tool performs the following tasks in this phase:

1. Performing a final SnapMirror update from the 7-Mode primary volume to the clustered Data ONTAP primary volume
2. Breaking and deleting the SnapMirror relationship between the 7-Mode primary volume to the clustered Data ONTAP primary volume
3. Enabling administrative access to the clustered Data ONTAP primary volume
4. Applying volume and services configuration to the Vserver
5. Setting up cluster peer and Vserver peer relationships between the primary and secondary clusters
6. Taking the 7-Mode volumes offline and removing the 7-Mode IP addresses selected for transition

7. Breaking and deleting the SnapMirror relationship between the 7-Mode primary volume and the clustered Data ONTAP secondary volumes
8. Configuring the data LIFs on the Vserver
9. Setting up a SnapMirror relationship between the primary and secondary clustered Data ONTAP volumes

The following illustration depicts the transition of a primary volume:



Resuming client access

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

1. Create the DNS entry for the domain names, which is generated from the domain name specified in the options `dns.domainname` command on the 7-Mode system, using the following command:

```
dns create -vserver vserver-name -domains <comma-separated-domain-names-list> -name-server comma-separated-name-servers-list
```

Example: If options `dns.domainname` on 7-Mode is set to `example.host1.com`, then the DNS entry should be configured on the cluster after the transition is complete using the following list of search domains.

```
dns create -vserver vs1 -domains example.host1.com -name-server  
192.5.10.10
```

2. Configure the features, such as IPv6 and FPolicy, that were available on the 7-Mode system, but were not transitioned automatically to the Vserver by the tool.
3. Ensure that the Vserver is ready to serve data to the clients by verifying the following:
 - The volumes on the Vserver are online and read/write.
 - The transitioned IP addresses are up and reachable on the Vserver.
4. Redirect client access to the clustered Data ONTAP volumes.

Related concepts

[How you transition a stand-alone volume](#) on page 33

[Features and volumes not supported for transition](#) on page 8

[Restrictions for transition](#) on page 14

Mode types you can use for completing the transition

You can complete the transition in either normal mode or test mode. Normal mode involves removing access to the 7-Mode volumes and IP addresses, completing data and configuration transition, and enabling administrative access to the clustered Data ONTAP volumes. In test mode, clients can continue to access the 7-Mode volumes.

Normal mode	The tool performs the following operations when transition is completed in normal mode: • Performs a final data copy from the 7-Mode volumes to the clustered Data ONTAP volumes • Establishes a SnapMirror relationship between the clustered Data ONTAP primary and secondary volumes if the session type is <i>primary</i> • Establishes a SnapMirror relationship between the 7-Mode primary volume and the clustered Data ONTAP secondary volume if the session type is <i>secondary</i> • Applies the 7-Mode configurations (including protocols and services configuration) to the clustered Data ONTAP volumes • Takes the 7-Mode volumes offline • Removes the 7-Mode IP addresses that are selected for transition • Configures the data LIFs on the Vserver with IP addresses • Sets the administrative state of the data LIFs on the Vserver to up • Enables administrative access to all the clustered Data ONTAP volumes
Test mode	In the test mode, the transitioning 7-Mode volumes are not taken offline and the transitioning IP addresses are not removed from the 7-Mode system. The tool performs the following operations when transition is completed in test mode:

- Performs a final data copy from the 7-Mode volumes to the clustered Data ONTAP volumes
- Establishes a SnapMirror relationship between the clustered Data ONTAP primary and secondary volumes if the session type is *primary*
- Establishes a SnapMirror relationship between the 7-Mode primary volume and clustered Data ONTAP secondary volume if the session type is *secondary*
- Applies the 7-Mode configurations (including protocols and services configuration) to the clustered Data ONTAP volumes
- Configures the data LIFs on the Vserver with IP addresses
- Sets the administrative state of the data LIFs on the Vserver to down

Related tasks


[*Completing a subproject*](#) on page 76

[*Completing the transition*](#) on page 139

States of a transition subproject or session

You should be aware of what the various states of a transition subproject or session are to manage the transition process. These states include preparation, copy baseline, copy update, copy pause, cutover, complete, aborted, and deleted.

The following image shows the different states in the transition process:



A transition subproject or session goes through the following states in a normal transition workflow:

- **Preparation**
When a transition subproject or session is created, it enters the preparation state.
In this state, 7-Mode and cluster information, such as host names, volume names, and IP addresses, is collected. You must run a precheck in this state and resolve all the errors, if any, to proceed with the transition process.
- **Copy baseline**
When a transition subproject or session is started, a SnapMirror relationship is established between the 7-Mode and clustered Data ONTAP volumes and the baseline copy is started.
- **Copy update**
A transition subproject or session moves to the copy-update state when the baseline transfer is completed for all the volumes.
In this state, the clustered Data ONTAP volumes are periodically updated according to the schedule specified in the subproject or session.
- **Cutover**
A transition subproject or session moves to the cutover state when you initiate the cutover for the volumes.
You must stop client access to the 7-Mode volumes before initiating cutover. A final update transfer is performed and the 7-Mode IP addresses are removed and configured on the Vserver. The SnapMirror relationship is broken and deleted. The duration of the cutover phase depends on the time taken for the final data transfer, for collecting all configurations from the 7-Mode system, and applying the configurations to the Vserver.
- **Complete**
The subproject or session moves to the completed state when cutover is successful and the 7-Mode configuration is applied to the clustered Data ONTAP volumes.
For a secondary subproject or session, a volume SnapMirror relationship is established with the 7-Mode primary volume. For a primary subproject or session, a volume SnapMirror relationship is established with the secondary cluster.

There are additional states that a subproject or session can go through depending on the transition operation:

- **Copy pause**
A transition subproject or session moves to the copy-pause state if you pause the subproject or session.
You can pause a subproject or session in any state. All SnapMirror transfers are paused in this state.
You can resume a paused subproject or session. If you paused the subproject or session when transition was in the copy-baseline state, the resumed subproject or session returns to the copy-baseline state. Similarly, if you paused the subproject or session at copy-update, the resumed subproject or session returns to the copy-update state.
- **Copy-unknown**
This is a transient state and a subproject or session might move to this state due to any of the following causes:

- Baseline transfer fails
- Tool is unable to retrieve the volume status
- Transition server is restarted and the tool is unable to retrieve the credentials of the cluster
- Invalid cluster credentials (for example, if the system password is changed during transition)

You can start a subproject when it is in the copy-unknown state; however, it fails if the previous state of the subproject was copy-update.

You can start a subproject successfully if the previous state of the subproject was copy-baseline.

- Aborted

You can abort a subproject or session when it is in the copy-baseline, copy-update, copy-pause, or cutover state.

When you abort a subproject or session, all the SnapMirror relationships between the 7-Mode system and the Vserver that are created during transition are deleted. After a subproject or session moves to the aborted state, the only operation possible on the subproject or session is to delete it.

Related concepts

[How transition works](#) on page 10

[Types of transition subprojects or sessions](#) on page 33

Considerations for using SnapMirror for transition

You must be aware of the different considerations for running transition operations along with existing 7-Mode to 7-Mode SnapMirror or SnapVault operations.

There are a limited number of replication streams available on each 7-Mode controller. The number of replication streams available varies by system model. For information about the maximum number of concurrent operations for your system model, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

The data copy operations for transition started by 7-Mode Transition Tool use replication streams on the 7-Mode controller. Therefore, 7-Mode Transition Tool must contend with both 7-Mode-to-7-Mode SnapMirror or SnapVault operations and 7-Mode-to-cluster mode SnapMirror operations for replication streams on the 7-Mode controller.

You can ensure that your existing backup and disaster recovery (DR) operations are not impacted by 7-Mode Transition Tool transition operations in the following ways:

- Ensure SnapMirror schedules selected for a session or subproject do not overlap with existing SnapMirror or SnapVault schedules
- Transition operations finish faster by using multiple data copy paths
- If sufficient replication streams are not available for 7-Mode-to-7-Mode scheduled incremental updates and the 7-Mode-to-7-Mode scheduled incremental updates require prioritization, make the required replication streams available on the 7-Mode controller by pausing one or more of the 7-Mode Transition Tool sessions that are using the same 7-Mode controller.

Attention: The tool runs a precheck when you initiate data copy on a subproject or a session. You are alerted by a precheck warning if sufficient replication streams are not available.

Ensure SnapMirror schedules do not overlap

Incremental SnapMirror updates from the 7-Mode controller to the cluster start at the schedule you provide while creating the session. You must provide schedules that do not overlap with any of the 7-Mode-to-7-Mode SnapMirror or SnapVault operations and ensure that the 7-Mode-to-7-Mode SnapMirror or SnapVault operations do not fail because the required number of replication streams are not available.

Start the baseline when enough replication streams are available on the 7-Mode controller. The baseline operation will use replication streams on the 7-Mode for the duration of the operation.

The 7-Mode-to-cluster mode incremental SnapMirror updates will also utilize replication streams on the 7-Mode controller for the entire duration of the update operation. If a baseline copy operation is happening for another session using the same 7-Mode controller, the 7-Mode to cluster mode scheduled incremental updates might fail due to insufficient replication streams.

The cutover operation will force a final incremental update for all the volumes in a session. Make sure that there are enough replication streams available on the 7-Mode controller for the final update to finish successfully. To achieve this, you might need to pause 7-Mode-to-7-Mode SnapMirror or SnapVault or 7-Mode-to-cluster mode SnapMirror operations for the duration of the cutover operation. The cutover operation will fail if the final incremental update operation fails even for one of the volumes in the session.

You must consider the following for secondary sessions:

- The secondary cutover operation will establish a 7-Mode-to cluster-mode, cross-mode SnapMirror relationship by performing a resync operation between the primary 7-Mode volumes and the secondary volumes on the Vserver. You must ensure that there are sufficient replication streams available on the 7-Mode primary controller for the resync to complete successfully.
- After cutover, the incremental SnapMirror updates for 7-Mode-to-cluster mode, cross-mode SnapMirror relationship will begin based on the schedules carried over from 7-Mode to 7-Mode SnapMirror schedules. You must ensure that there are sufficient replication streams available on the 7-Mode primary controller.

Note: If you performed the transition in test mode, these schedules will overlap with the existing 7-Mode-to-7-Mode schedules.

Using multiple paths for transition

The 7-Mode Transition Tool 1.2 enables you to specify two paths for transition by using a data copy IP address and a multipath IP address. However, both paths can be used only for load-balancing, not for failover.

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

- You must have two intercluster LIFs on each node of the cluster to use multiple paths for transition.
- If the secondary path fails before the transfer is started, the baseline transfer begins after a delay.
- If both paths are available when the data transfer operation between the 7-Mode controller and the cluster starts, both paths must remain available throughout the data transfer. If any path fails during data transfer, the data transfer operation fails.
- After the transfer starts between the 7-Mode system and the clustered Data ONTAP system, both paths must be available for SnapMirror transfers to continue. If any path fails when the transfer is in progress, the data copy fails.

Related concepts

[*Guidelines for deciding when to perform cutover*](#) on page 47

Related references

[*Limits for transition*](#) on page 48

Related information

[*Documentation on the NetApp Support Site: support.netapp.com*](#)

Guidelines for deciding when to perform cutover

Because transition cutover is disruptive to clients, you must plan the activity to minimize the downtime. You must schedule the cutover during a low-activity window. Also, you must monitor the transfer duration for the last few update transfers and confirm that the final copy-update does not take a long time.

You must keep monitoring the SnapMirror status for each volume. If the last transfer duration of the previous few updates for the volume is within an acceptable limit, most of the data changes in the volume should have been copied and the time for final data update during cutover should be within the acceptable limit.

You can derive the approximate downtime depending on the number of volumes that are transitioned. Applying configurations on the clustered Data ONTAP volumes consumes additional time.

To minimize the cutover time, the network latency between the 7-Mode Transition Tool and storage systems should be minimum. For transitioning a volume SnapMirror relationship, the network latency between the tool and the primary systems should be minimum.

Related concepts

[*Considerations for using SnapMirror for transition*](#) on page 45

Related references

[Limits for transition](#) on page 48

Limits for transition

When you transition volumes using the 7-Mode Transition Tool, you should consider certain limits for transition such as the number of volumes in a session or subproject, and the number of sessions or subprojects that can run simultaneously.

Parameter	Maximum limit
Number of volumes in a subproject or session (7-Mode Transition Tool web interface and CLI)	20
Maximum number of 7-Mode controllers from which volumes can be transitioned at any given time (7-Mode Transition Tool web interface only)	Two 7-Mode controllers, which can be one pair of High Availability (HA) controllers or two non-HA controllers. For this limit, only primary volumes and stand-alone volumes are considered; controllers that hold secondary volumes included in projects are not counted towards this limit. The 7-Mode Transition Tool web interface does not allow you to transition stand-alone or primary volumes that involve a third 7-Mode controller. This limit does not apply for 7-Mode Transition Tool CLI.
Maximum number of clusters to which volumes can be transitioned at any given time (7-Mode Transition Tool web interface only)	Two clusters for stand-alone and primary volumes involved in the transition. The clusters used to transition secondary volumes with the SnapMirror relationship are not counted. This limit does not apply for 7-Mode Transition Tool CLI.
Maximum number of active projects that can be managed at any given time (7-Mode Transition Tool web interface only)	Eight active projects (each project can have either a stand-alone subproject or a primary and a secondary subproject.) Note: The tool supports listing of more than 100 projects; however, it is not recommended. This limit does not apply for 7-Mode Transition Tool CLI.

Important: The restriction of two 7-Mode controllers and two clusters in the 7-Mode Transition Tool web interface is applicable across subprojects created in the web interface and sessions created in the CLI. Therefore, if you create a session in the CLI that uses one 7-Mode controller and transitions volumes to one cluster, you can create projects in the web interface that uses the 7-Mode controller and cluster that is used in the CLI session and one additional 7-Mode controller and one additional cluster.

Related concepts

[*Considerations for using SnapMirror for transition*](#) on page 45

[*Guidelines for deciding when to perform cutover*](#) on page 47

What the transition configuration file is

The `\etc\conf\transition-tool.conf` file is installed with the 7-Mode Transition Tool and contains information about all the configurable options that are necessary for the tool to operate properly.

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

You can edit the file to modify any configuration option that is used by the tool. After you modify any option, you must restart the NetApp Transition Tool service for the new values to take effect. After the NetApp Transition Tool service is restarted, you must add the 7-Mode system and cluster credentials again.

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

Related references

[*Limits for transition*](#) on page 48

Severity levels for precheck messages

You can validate the feasibility of transition and verify if the 7-Mode volumes can be transitioned by running the transition precheck operation. Transition precheck reports all transition issues. Transition

issues are assigned different severity levels, depending on the impact of the issue on the transition process.

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

- Initial prechecks: You can run the initial prechecks when you have only selected the 7-Mode volumes for transition.
- Final prechecks: You can run the final prechecks on the 7-Mode volumes selected for transition and the Vservers and aggregates to which the selected volumes will be transitioned.

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

Error	Configurations that cannot be transitioned. You cannot continue the transition if there is even one error. The following are some example configurations on the 7-Mode system that cause an error: • Traditional volumes • SnapLock volumes • LUNs • Volumes are offline
Warning	Configurations that can cause minor problems after transition. You can continue the transition with these warnings. However, after the transition you might lose some of these configurations or might need to complete some manual tasks for enabling these configurations in clustered Data ONTAP. The following are some example configurations on the 7-Mode system that generate a warning: • MetroCluster configuration • IPv6 • NFSv2 • NDMP configurations • Interface groups and VLANs • Routing Information Protocol (RIP)
Information	Configurations that have been successfully transitioned.

Related tasks

[Assessing controllers for transition](#) on page 71

[Running prechecks](#) on page 77

[Running precheck on a session](#) on page 135

Considerations for transition

You must be aware of considerations about CIFS, NFS, and local users and groups before running the transition operations.

Considerations for NFS

Before running the transition operations, you should be aware of some important considerations for NFS transition.

Considerations for NFS exports transition

The behavior of NFS exports in clustered Data ONTAP differs from Data ONTAP operating in 7-Mode because there are no junction paths in 7-Mode. In clustered Data ONTAP, to mount a volume or a qtree the NFS client must have read-only permissions at all the parent junction paths up to the Vserver's root volume junction path (/).

Transition tool exports the Vserver root volume to allow read-only access for all the NFS clients

To ensure that all the transitioned volumes or qtrees are mountable, the 7-Mode Transition Tool exports the root volume of the Vserver to allow read-only access for all the NFS clients.

As part of the cutover operation, the tool checks if the Vserver root volume is exported to allow read-only access to all NFS clients. If the Vserver root volume is not exported to allow read-only access to all NFS clients, then the 7-Mode Transition Tool creates a new export policy to allow read-only access for all the NFS clients and exports the root volume of the Vserver with the newly created export policy.

Transition of 7-Mode volumes or qtrees with exports having configurations that are not supported on clustered Data ONTAP

When 7-Mode volumes that have export configurations which are not supported on clustered Data ONTAP are transitioned, the 7-Mode Transition Tool exports these volumes to allow read-only permissions to all NFS clients on the Vserver.

When the qtrees that have export configurations, which are not supported on clustered Data ONTAP are transitioned, they inherit the export policy of the parent volume.

Post-transition manual steps

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

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

Behavioral difference in qtree exports between Data ONTAP operating in 7-Mode and clustered Data ONTAP

Use case 1

In clustered Data ONTAP, for an NFS client to mount a qtree, that NFS client must have read-only permissions at all the parent junction paths up to the Vserver's root volume junction path (that is, /). This behavior in NFS exports is different from Data ONTAP running 7-Mode because there are no junction paths in 7-Mode.

Post-transition manual steps

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

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

The following is an example in the 7-Mode storage system (192.168.26.18):

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

On the 7-Mode storage system, the NFS client 192.168.10.10 can mount to the qtree export path.

```
[root@192.168.10.10 ]# mount 192.168.26.18:/vol/volstd10/qtree1
transition_volume_qtree
[root@192.168.10.10 ]#
```

The following is an example of an export policy of the volume volsdt10 on clustered Data ONTAP after the transition:

```
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
```

```

cluster-01::> Allow Creation of Devices: true

```

The following is an example of an export policy of the qtree qtree1 on the clustered Data ONTAP after the transition:

```

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

Vserver: std_22
Policy Name: std_2225
Rule Index: 1
Access Protocol: any
Client Match Hostname, IP Address, Netgroup, or Domain: 192.168.10.10
RO Access Rule: sys
RW Access Rule: never
User ID To Which Anonymous Users Are Mapped: 65534
Superuser Security Types: none
Honor SetUID Bits in SETATTR: false
Allow Creation of Devices: true
cluster-01::>

```

After the transition is complete, the NFS client 192.168.10.10 cannot mount the transitioned qtree. 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 is an example of clustered Data ONTAP (Vserver data LIF 192.168.35.223):

```

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

```

Required manual step:

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

```

cluster-01::> export-policy rule create -vserver std_21 -policyname
std_2124 -clientmatch
192.168.10.10 -rorule sys -rwrule never -allow-suid false -allow-dev
true -superuser none -protocol nfs
(vserver export-policy rule create)

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

Vserver: std_22
Policy Name: std_2226
Rule Index: 1

```

```

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

Vserver: std_22
Policy Name: std_2226
Rule Index: 1
Access Protocol: nfs
Client Match Hostname, IP Address, Netgroup, or Domain: 192.168.10.10
  RO Access Rule: sys
  RW Access Rule: never
User ID To Which Anonymous Users Are Mapped: 65534
  Superuser Security Types: none
  Honor SetUID Bits in SETATTR: false
  Allow Creation of Devices: true

cluster-01:>

```

Use case 2

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

The following is an example of an 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]#

```

Similarly, in clustered Data ONTAP, the NFS client 192.168.10.10 has only read-only access to the qtree “qtree1” when the client accesses the qtree directly or through the mount point of the qtree's parent volume.

Post-transition manual steps

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

NFS only configuration transition

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

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

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

Considerations for CIFS

You must consider some important considerations for CIFS transition before running the transition operations.

Considerations for transitioning CIFS local users and groups

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

- If the destination cluster is running clustered Data ONTAP 8.2, transition should not be attempted for 7-Mode volumes that are serving CIFS data and are being accessed by local users and groups because 7-Mode Transition Tool does not support the transition of local users and groups to clustered Data ONTAP 8.2.
- 7-Mode Transition Tool supports the transition of local users and groups to clustered Data ONTAP 8.2.1.

However, transition of CIFS data-serving volumes from a 7-Mode controller or a vFiler unit that has local users and groups to a Vserver that has non-BUILTIN CIFS local users and groups is not supported. The Vserver must have only BUILTIN CIFS local users and groups for such transitions to proceed.

While transitioning local users and groups from a given 7-Mode controller or a vFiler unit to a given Vserver, local users and groups from the first transition subproject or session are transitioned. In the subsequent transition of subprojects or sessions with the same 7-Mode controller or vFiler unit to the same Vserver, the transition of local users and groups is ignored; although, transition succeeds. The local user's name on the 7-Mode must not be the same as the CIFS server name on the Vserver.

- You must be aware of the limits on the number of local users and groups supported in clustered Data ONTAP 8.2.1.

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

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

After the transition is complete, you can access these users from the Windows system by using the password `cifsUser@1`. You must then manually change the password for such CIFS local users on the Vserver using the following clustershell CLI command: `cifs users-and-groups local-user set-password -vserver <vserver name> -user-name <user-name>`

Transitioning volumes using the 7-Mode Transition Tool (web interface)

The 7-Mode Transition Tool (web interface) is a web-based application that enables you to transition 7-Mode volumes to clustered Data ONTAP.

You configure, execute, and manage a transition process in the web interface by using projects and subprojects. Projects are the containers for subprojects. The actual configuration, execution, and management of the transition is done by using subprojects.

Related concepts

[Transitioning from 7-Mode to clustered Data ONTAP](#) on page 7

[Preparing for transition](#) on page 12

[Understanding the 7-Mode Transition Tool](#) on page 32

[Understanding the transition process](#) on page 57

Understanding the transition process

The transition process involves creating a transition project and running it. You use the same interface for transitioning stand-alone volumes and volumes with SnapMirror relationships. However, the workflows for transitioning stand-alone volumes and volumes with SnapMirror relationships are different.

The transition process involves the following two phases:

- Creating a project using the New Transition Project wizard
This includes grouping similar volumes into subprojects.
- Starting and completing a project by using Dashboard

In each phase, the workflow for transitioning volumes as stand-alone volumes is different from the workflow for transitioning volumes with SnapMirror relationships.

Related concepts

[How transition works](#) on page 10

[States of a transition subproject or session](#) on page 42

[Features and volumes not supported for transition](#) on page 8

[Restrictions for transition](#) on page 14

Related references

[Limits for transition](#) on page 48

Project creation workflows

When you create a project by using the New Transition Project wizard, you first group volumes into subprojects and then decide how these groups of volumes will transition by configuring the subprojects. While the process of configuring individual subprojects is the same, the workflow for creating a project varies based on the types of subprojects in that project.

The key features of the project creation workflows are:

- You can add the controllers and clusters that you need for the project.
- A number of restrictions are enforced to ensure that only those volumes that can be transitioned together are grouped in a subproject.
- You can run prechecks at the following three stages in the wizard:
 - After you add the 7-Mode volumes that you want to transition
 - After you map the volumes to aggregates on the cluster
 - After you have selected SM schedule for incremental transfers


Related concepts

[*Types of transition subprojects or sessions*](#) on page 33

Steps to transition stand-alone volumes

You need to understand how to create a project that transitions stand-alone volumes if you want to separately transition a group of volumes as stand-alone volumes. You can add volumes from the vFiler unit to a stand-alone project.

The following illustration shows the steps for creating a project that transitions stand-alone volumes:



Related concepts

[Types of transition subprojects or sessions](#) on page 33

[Features and volumes not supported for transition](#) on page 8

[Restrictions for transition](#) on page 14

[Restrictions on creating and configuring projects](#) on page 66

[Requirements and restrictions for using the web interface](#) on page 66


Related references

[Limits for transition](#) on page 48

How you create a project that transitions SnapMirror relationships

You need to understand how to create a project that transitions volumes with SnapMirror relationships to plan the transition.

The following illustration shows the steps involved in creating a project that transitions volumes with SnapMirror relationships:



Projects that transition SnapMirror relationships have two subprojects that can be configured, started, and completed separately. The workflow for configuring both the subprojects is the same. You use the New Transition Project wizard to complete the configuration of the first subproject and save it before you start configuring the second one.

The process for creating a project that transitions volumes with SnapMirror relationships is different from the process for creating a project that transitions stand-alone volumes in the following ways:

- In the Volume selection pane, you use the Transition with SnapMirror relationship column, instead of the Transition as Stand-alone column to add primary or secondary volumes.
- You have to configure two subprojects, instead of one for stand-alone projects.

Related concepts

[Types of transition subprojects or sessions](#) on page 33

[Features and volumes not supported for transition](#) on page 8

[Restrictions for transition](#) on page 14

[Restrictions on creating and configuring projects](#) on page 66

[Requirements and restrictions for using the web interface](#) on page 66

Related references

[Limits for transition](#) on page 48

Understanding the project start and completion workflows

You need to understand the start and completion workflows to manage the data transfer and cutover operations.

Related concepts

[Types of transition subprojects or sessions](#) on page 33

[Mode types you can use for completing the transition](#) on page 41

[Guidelines for deciding when to perform cutover](#) on page 47

Dashboard

You can edit subprojects, run prechecks and review the results, and start and complete subprojects from Dashboard. You can also review the status of the running subprojects and delete projects that have no active subprojects. In addition, you can view the transition history of completed and aborted projects.

Following are the key controls on the Dashboard page:

Project navigation The tree control on the left of the Dashboard lists all the projects (and the subprojects within them) that you have created. It also allows you to view the transition history of completed and aborted projects. The transition history is grouped to display the projects in the following order - active projects, completed projects, and aborted projects. The projects are color-coded. For more information, see the legend provided at the bottom of the pane.

Delete project	You can delete projects that have either completed or have not started. If you want to delete a project that has been started, you need to first abort the subproject that has been started.
Start	You can select a subproject and click this button to start the transition for the selected subproject.
Cutover	You can select a subproject and click this button to complete the selected subproject.
Target volume status	This table displays the status of data copy operations between the 7-Mode volumes and the target volumes for the selected subproject.
View transition details	You can select a subproject and click this button to view the detailed status of the target volumes for the selected subproject.
Refresh	You can select a subproject and click this button to refresh the Target Volume Status table. The Target Volume Status table is refreshed automatically if the auto-refresh option is selected.
Network address to be migrated	This table displays the network address or addresses to be migrated for the selected subproject.
Edit or remove LIF	You can edit or remove any of the LIFs listed in the Network Address To Be Migrated table.
Add LIF	You can add a LIF to the selected subproject using this button.
Run prechecks	You can select a subproject and click this button to run prechecks on the selected subproject.
See precheck results	You can select a subproject and click this button to view detailed precheck results in a dialog box. You can also save the results to a CSV (comma-separated values) file using the Save as CSV option in the Precheck Results dialog box.
Edit configuration	You can select a subproject and click this button to edit that subproject.
Pause	You can select a subproject that has been started and click this button to pause the data copy operations.
Resume	You can select a subproject that has been paused and click this button to resume the data copy operations.
Abort	You can select a subproject that has been started and click this button to abort the data copy operations for the selected subproject.
Auto refresh	You can select this option to enable automatic refresh of the data displayed in the View transition details table. The auto refresh interval in the web interface is five minutes.

Related concepts

[Severity levels for precheck messages](#) on page 49


[How you start and complete a project that transitions stand-alone volumes](#) on page 63

[How you start and complete a project that transitions volumes with SnapMirror relationships](#) on page 64

How you start and complete a project that transitions stand-alone volumes

You need to understand how to start and complete projects that transition stand-alone volumes to manage and complete such transitions.

The following illustration shows the steps involved in starting and completing a project that transitions stand-alone volumes:



You can manage this workflow in the following ways:

- You can pause and resume a subproject any time after it has started. However, you cannot edit the configuration if after you have started it.
- You can abort a subproject that you have started and then delete the project that contains it.
- After the baseline transfer is complete for all the volumes in the subproject, the tool performs incremental data transfers at the intervals specified in the selected SnapMirror schedule that you selected while configuring the subproject.

At this stage, you can initiate cutover to complete the transition.

Related concepts

[*Severity levels for precheck messages*](#) on page 49


[*Mode types you can use for completing the transition*](#) on page 41

[*Guidelines for deciding when to perform cutover*](#) on page 47

How you start and complete a project that transitions volumes with SnapMirror relationships

You need to understand how to start and complete a project that transitions volumes with SnapMirror relationships and the various states it goes through, to complete the transition. The transition for volumes with SnapMirror relationships is complete when both the primary and secondary subprojects are complete.

The following illustration shows the steps for starting and completing a project that transitions volumes with SnapMirror relationships:



You must complete the secondary subproject before you complete the primary subproject. This ensures that the SnapMirror relationship is maintained after the transition to clustered Data ONTAP. When the transition of only the secondary volumes is complete, a SnapMirror relationship is formed between the 7-Mode primary volume and the secondary volume created on the cluster.

If you choose to transition only the secondary volumes, you do not need to start or complete the primary subproject. In this case, the SnapMirror relationships are between the primary volumes on the 7-Mode and the secondary volumes on clustered Data ONTAP.

Related concepts

[Severity levels for precheck messages](#) on page 49

[Mode types you can use for completing the transition](#) on page 41

[Guidelines for deciding when to perform cutover](#) on page 47

Requirements and restrictions for using the web interface

You must be aware of the requirements and restrictions of using the 7-Mode Transition Tool (web interface).

Requirements of the web interface

- The NetApp Transition Tool service must be running on the machine on which it was installed. The service is set to automatic by default, and should start when you restart the machine.

Restrictions of the web interface

- You should not use the web interface and the command-line interface (CLI) simultaneously or interchangeably.
For example, if you create a project using the web interface, you should start and complete the project using the web interface. Similarly, if you define a transition session using the CLI, you should perform all the transition operations on that session using the CLI.
- The tool does not support web browser navigation. If you use the Back and Forward browser buttons, you might see unwanted results. If you use the Refresh browser button, you will be directed to the Login page.

Related concepts

[Restrictions for transition](#) on page 14

[Limits for transition](#) on page 48

[Restrictions on creating and configuring projects](#) on page 66

[Managing a transition project](#) on page 77

Related tasks

[Creating and running a transition project](#) on page 69

Restrictions on creating and configuring projects

You need to know the restrictions on creating and configuring projects to understand why in the New Transition Project wizard, some volumes and vFiler units are disabled on the Volume selection pane and why some Vservers and aggregates are disabled on the Volume mapping pane.

Restrictions on controllers and vFiler units

- You cannot create a project if the new project uses a controller that is not one of the two controllers used by the existing active projects.

- The Volume selection pane allows you to add volumes from a single vFiler unit only because a subproject can include volumes from a single vFiler unit.
When creating the project, the volumes included in the project are removed if you select a different vFiler unit.

Restrictions on clusters, Vservers, and aggregates

- You can use only two clusters at any given time across all the existing projects.
Therefore, you cannot select a cluster if it is not one of the two clusters used in existing active projects. This restriction applies to stand-alone volumes and primary volumes selected for transition with their SnapMirror relationships. The number of clusters used by secondary volumes that have been selected for transition with their SnapMirror relationships is not considered.
- You cannot include a Vserver if the Vserver is part of an active subproject (a subproject that has not been completed or aborted) that transitions volumes from a different vFiler unit.
- Aggregates that are in the failed or offline state are not displayed.

Restrictions on selecting volumes for transition

Volumes are grouped into subprojects primarily on the basis of the type of transition that needs to be performed. Therefore, the following restrictions apply on selecting volumes:

- In a stand-alone subproject, you can include only those volumes that you want to transition as stand-alone volumes. You can transition stand-alone, primary, and secondary volumes as stand-alone volumes.
- In a primary subproject, you can include only those primary volumes that you want to transition with their SnapMirror relationships.
Secondary volumes are automatically selected when you select the associated primary volumes for transition with their SnapMirror relationships. You can also include any volume that you want to transition as a stand-alone volume in either the primary subproject or the secondary subproject.

Restrictions on selecting volumes for transition with SnapMirror relationships

- To transition volumes with their SnapMirror relationships, you must add the primary volume to the subproject using the **Transition with SnapMirror relationship** option.
You cannot use secondary volumes to select SnapMirror relationships for transition.
- After you include the first primary volume that has a SnapMirror relationship for transition (thereby automatically including the associated secondary volume), you can include only those primary volumes that have secondary volumes on the same vFiler unit as the associated secondary volume.
You cannot select a primary volume for transition that has a SnapMirror relationship if its secondary volume is on a different vFiler unit (and not the vFiler unit that contains the secondary volume of the first primary volume included in the project.)
- For a new project, you cannot select volumes that are already part of active subprojects.
- You cannot select a volume that has a SnapMirror relationship for transition if that volume is source (primary) in one SnapMirror relationship and destination (secondary) in another.

- You cannot select a volume for transition that has a SnapMirror relationship if that volume has more than one secondary volume.

Note: You can transition such volumes with their SnapMirror relationships by using the 7-Mode Transition Tool's command line interface. In the web interface, you can transition such volumes as stand-alone volumes only.

Related concepts

[Features and volumes not supported for transition](#) on page 8

[Restrictions for transition](#) on page 14

Related references

[Limits for transition](#) on page 48

Prerequisites

You must be aware of the prerequisites at each stage of the different workflows in running a transition project.

Common prerequisites for all stages

- You must ensure that the 7-Mode controllers and clusters that you want to include in the transition are reachable from the Windows host where the tool is installed.
- You must have all administrator-level privileges for the controllers and clusters that you want to include in the transition project.

Prerequisites for creating projects

- You must have created the required Vservers on the cluster.
- You must have identified or created a SnapMirror schedule on the cluster.
This SnapMirror schedule is required for incremental data transfer during transition.
- You must have saved and closed any project that is open for editing.

Prerequisites for starting data copy for a subproject

- You must have created and configured a transition project.
- You must have ensured that you have addressed the issues identified in the precheck results.

Prerequisites for performing cutover for a subproject

- You must have ensured that the baseline copy has completed successfully.
- For projects that transition volume SnapMirror relationships, you must have completed the secondary subproject before completing the primary subproject.

- You must have ensured that you have prepared for a short downtime of the 7-Mode controller from which volumes are being transitioned.

Logging in to the tool

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

Steps

1. Navigate to the login page.

If you are accessing the web application...	Then...
From the server on which it was installed	Double-click the 7-Mode Transition Tool icon on your desktop.
From another computer	Navigate to <code>https://IP_address:port/transition</code> or <code>http://IP_address:port/transition</code> by using one of the supported web browsers. In this URL, <i>IP_address</i> is the IP address of the server and the <i>port</i> could be 8443 (default), the value specified for the <i>tool.https.port</i>, or 8088, the value specified for the <i>tool.http.port</i> parameter in the <i>transition-tool.conf</i> file.

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

Note: If you are using the 8088 port (`http://IP_address:port/transition`), then the self-signed SSL certificate message is not displayed.
3. Log in to the application using your Windows credentials for the server.

Creating and running a transition project

To transition volumes using the 7-Mode Transition Tool (web interface), you must first create a project and configure its subprojects, and then start and complete all the subprojects created in that project.

Before you begin

- You must have reviewed the following information:

[Prerequisites](#) on page 68

[Requirements and restrictions on the project workflow](#) on page 66

[Guidelines for selecting and grouping volumes](#) on page 66

Steps

1. [Adding controllers and clusters](#) on page 70
2. [Assessing controllers for transition](#) on page 71
3. [Creating a transition project](#) on page 72
4. [Starting a subproject](#) on page 75
5. [Completing a subproject](#) on page 76

Related concepts

[Project creation workflows](#) on page 58

[Understanding the project start and completion workflows](#) on page 61

[Managing a transition project](#) on page 77

[Troubleshooting web interface issues](#) on page 83

Adding controllers and clusters

You can add the controllers and clusters that are required for the transition by using the Storage Systems page. You can also remove the controllers and clusters using the Storage Systems page.

Before you begin

You must have ensured that you meet the generic and specific prerequisites for adding controllers and clusters listed in [Prerequisites](#) on page 68.

About this task

Although you can add and remove controllers and clusters by using the Storage Systems page, the tool does not allow you to remove the controllers and clusters that are part of any active subproject or session.

The 7-Mode controller and cluster information that you provide on this page is not persistent. If the NetApp Transition Tool service is restarted for any reason, the tool will prompt you for information about controllers and clusters that are part of active projects. Active projects are projects that have not completed or aborted.

Note: If a 7-Mode controller is part of a high availability (HA) pair, the tool does not ask for credentials of the HA partner of the 7-Mode controller (unless the HA partner is part of another active subproject.)

The Storage Systems page also enables you to assess the transition readiness of up to 20 controllers. The multi-controller assessment (precheck) feature assesses the controllers and all the volumes on them.

Steps

1. Click **Storage Systems**.

2. In **Hostname**, enter the fully qualified domain name (FQDN) or the IP address of the 7-Mode controller or the clustered Data ONTAP system.
For a cluster, you can specify the IP address (or FQDN) of the management interface, or the IP address (or FQDN) of any node within the cluster. For a 7-Mode controller, you need to specify the IP address of the default vFiler unit, because the IP addresses of individual vFiler units are not accepted.
3. In **Username** and **Password**, enter the administrator credentials for the specified host, and click **Add**.
The tool detects the type of system added, and includes it in the correct column of the table.
4. Repeat Steps 2 and 3 to add all the controllers and clusters that you require for the transition.

Related concepts

[Features and volumes not supported for transition](#) on page 8

[Restrictions for transition](#) on page 14

[Preparing for transition](#) on page 12

[Understanding the 7-Mode Transition Tool](#) on page 32

Assessing controllers for transition

You can assess if the 7-Mode controllers you have included in the tool for transition meet the key requirements for transition. This procedure enables you to check the 7-Mode controllers (and all the vFiler units and volumes on them) for most of the common issues that prevent a successful transition.

Before you begin

- The 7-Mode controllers must be added as explained in [Adding controllers and clusters](#) on page 70.
- Precheck operation must not be running in any of the subprojects.
- The vFiler units on the 7-Mode controllers must either be running or stopped with the root volume in online state.

About this task

The multi-controller assessment (precheck) feature provides a single consolidated report that lists issues detected at the controller, vFiler, and volume levels. You can assess a maximum of 20 7-Mode controllers at a time using this feature.

Warning: Controller assessment fails if any controller that you select for assessment is not accessible.

Steps

1. Click **Storage Systems** on the Home page.
2. Select the controllers you want to check from the Assess Controller column.

3. Click **Assess Controller(s)**.
4. Click **View results** from the **Controller assessment completed successfully** dialog box.

Issues detected are categorized in the same three categories that are used in the Precheck Results: **Errors**, **Warnings**, and **Information**.

The results page is displayed in a new browser tab. To view the results later, you must save the HTML file.

You can also save the raw results in CSV (comma separated values) format by clicking on the **csv file** link.

Related concepts

[Severity levels for precheck messages](#) on page 49

Related tasks

[Running prechecks](#) on page 77

[Running precheck on a session](#) on page 135

Creating a transition project

Creating a transition project includes selecting and mapping 7-Mode volumes to clustered Data ONTAP Vservers, mapping interfaces, running prechecks, and defining SnapMirror schedules for copying incremental data.

Before you begin

You must have reviewed and confirmed that you meet the generic and specific prerequisites for creating a transition project listed in [Prerequisites](#) on page 68.

Steps

1. On the **Home** page, click **New Transition Project**.
2. Optional: If the controller and cluster required by the new project are not listed, add them from the **Enter Device Credentials** pane, and click **Next**.
3. In the **New Transition Project Name** dialog box, enter a name for the project and click **Continue**.

Note: A project name can contain up to 25 alphanumeric characters. Special characters and spaces are not allowed.

4. From the **Choose Origin Volumes to Transition** section, select the 7-Mode controller or the vFiler unit from which you want to add volumes.
5. Add the volumes you want to include in the project by selecting one of the following:

If you want to transition the...	Then...
Volume as a stand-alone volume	Select Transition as stand-alone. A stand-alone subproject is created if you select the first volume from this column.
Primary volume and its corresponding secondary volume, while maintaining their SnapMirror relationship	a. Select Transition with SnapMirror Relationship for a primary volume. Two subprojects are created—a primary subproject for primary volumes and a secondary subproject for secondary volumes—if you select the first volume from this column. The selected volume and the corresponding SnapMirror primary or secondary volume are added to the appropriate subproject. b. If the controller for the corresponding SnapMirror volume is not included in the project, in the Need additional storage system credentials dialog box, enter the details for the controller to be included.

The volumes that cannot be added are grayed out in the Choose Origin Volumes to Transition pane.

Note: The transition tool supports transition of qtree level exports for Data ONTAP 8.2.1 and later clustered releases.

The hyperlink provided on the volume name, opens a dialog box that lists related qtrees and their attributes, such as Qtree, Owning vfiler, State, Security, NFS Export path, and NFS Actual path.

- Volumes with qtrees that act as secondary volume in qtree SnapMirror relationships (QSM) will continue to be unavailable for selection.
- Volumes with a qtree that acts as both source and destination for QSM will continue to be unavailable for selection.

6. After you have selected all volumes you want to include in the project, click **Create Project and Continue**.
7. Optional: From the **IP Configuration** pane, enter an additional IP address (in addition to the data copy IP) if you want to use multipath data copy and click **Next**.
8. Optional: From the **Prechecks** pane, click **Run Prechecks**.

The initial prechecks run on the 7-Mode volumes selected for transition.

The prechecks might take several minutes to complete. After the prechecks are complete, the Precheck page is displayed with the results.

9. To save the Precheck results, click **Save as CSV**.
10. Click **Next**.
11. From the **Vserver Mapping** pane, select the Vserver to which the 7-Mode volumes will be transitioned, and click **Next**.

If you select a Vserver that belongs to a clustered Data ONTAP 8.2 system, the **Do you have any Local users and groups accessing cifs shares or files configured on any of the below 7-mode storage system(s)?** dialog box is displayed.

12. Click **Yes** or **No** based on whether Local users and groups rules are configured.

Note: If you select **Yes**, a dialog box stating that local users and groups are not supported for migration is displayed, and you are asked to choose a cluster running clustered Data ONTAP 8.2.1. The 7-Mode Transition Tool transitions the 'Local users and Group' configurations (if any) automatically, if the target cluster is running clustered Data ONTAP 8.2.1.

If you select **No**, you can continue with the project creation.

13. The tool supports the transition of audit configuration from 7-Mode storage system to clustered Data ONTAP. To facilitate the transition of this configuration, provide a path on the destination Vserver by entering appropriate value in the **Vserver audit logs destination path** dialog box, which can be used for saving the audit logs.
14. From the **Volume Mapping** pane, select the aggregates to which the 7-Mode volumes will be copied.
15. Optional: To change the name a volume will have on the cluster, enter a different name in **Target Volume**.
16. Click **Next**.
17. Optional: From the **Interface Mapping** pane, provide information about the logical interfaces that will be created on the Vserver.

If you want to...	Then...
Transition an existing 7-Mode IP address	a. Select the 7-Mode IP address. b. Click Choose Node and Port. c. From the Choose Node and Port dialog box, select the target node and port to which the LIF will be migrated and click Save.
Create a new LIF	a. Click New LIF. b. In the Provide details for additional LIF in the destination cluster dialog box, enter the details for the new LIF and click Add.

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

18. Click **Next**.
19. From the **SnapMirror Schedule** pane, select the required SnapMirror schedule for transition, and click **Next**.
Cron jobs that start with sched_ and CronJob_ are not included in the list of SnapMirror schedules displayed.
20. Optional: From the **Final Precheck** pane, click **Run Prechecks**.

The Final precheck pane is similar to the Precheck pane described in Step 8. Review the precheck results, and click **Next**.

21. From the **Subproject Summary** pane, save the project by clicking one of the following:

If you want to...	Then...
Save the subproject and go to the Dashboard page	Click Save & Goto Dashboard .
Save the current subproject and configure the next subproject	Click Save & Configure Next Subproject . This button is displayed only for projects that have two subprojects.

Related concepts

[Project creation workflows](#) on page 58

[Restrictions on creating and configuring projects](#) on page 66

[Severity levels for precheck messages](#) on page 49

[Managing a transition project](#) on page 77

Related tasks

[Starting a subproject](#) on page 75

[Completing a subproject](#) on page 76

Related references

[Limits for transition](#) on page 48

Starting a subproject

After you create a project, you can start copying data from the 7-Mode volumes. You start copying data by starting individual subprojects.

Before you begin

You must have reviewed and confirmed that you meet the generic prerequisites and the specific prerequisites for starting a subproject listed in [Prerequisites](#) on page 68.

Steps

1. From the **Dashboard** page, select the subproject that you want to start.

If the precheck was not run or if it resulted in an error, the **Start** button is not available. If the precheck detects errors, you must perform the corrective actions specified in the precheck message and run the prechecks again.

2. Click **Start**.

The precheck is run once again in the background, and if no errors are detected, a baseline transfer is started. If the precheck detects errors, they are displayed in a dialog box.

3. Click **Refresh** to refresh the Target Volume Status table for the selected subproject.

The Target Volume Status table is updated with the latest status and progress values. The details shown in the table (and the details displayed when the **View Transition Details** button is clicked) are details of the SnapMirror relationships between the volumes in the transition. These details are not shown after the transition is completed or aborted.

Related concepts

[Requirements and restrictions for using the web interface](#) on page 66

[Severity levels for precheck messages](#) on page 49

[Managing a transition project](#) on page 77

Related tasks

[Creating a transition project](#) on page 72

[Completing a subproject](#) on page 76

Completing a subproject

You can complete a transition by completing the individual subprojects in the transition project.

Before you begin

You must have ensured that you meet the relevant prerequisites listed in [Prerequisites](#) on page 68.

Steps

1. From **Dashboard**, select the subproject that you want to complete.
2. Click **Cutover**.
A message is displayed warning you that the cutover operation is disruptive.
3. Click **Yes** to continue.
A message confirms that the cutover is successful.
Note: The cutover operation might take several minutes to complete.
4. Click **OK**.
Your data is now successfully transitioned to clustered Data ONTAP.

After you finish

If you have created a new logical interface (LIF), users and applications of the transitioned volumes must remap the drives by using the new IP addresses and ports after all the subprojects have been completed.

Related concepts

[Severity levels for precheck messages](#) on page 49

[Guidelines for deciding when to perform cutover](#) on page 47

[Managing a transition project](#) on page 77

[Requirements and restrictions for using the web interface](#) on page 66

Related tasks

[Creating a transition project](#) on page 72

[Starting a subproject](#) on page 75

Managing a transition project

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

Related concepts

[Understanding the project start and completion workflows](#) on page 61

[Severity levels for precheck messages](#) on page 49

[Requirements and restrictions for using the web interface](#) on page 66

[Guidelines for deciding when to perform cutover](#) on page 47

[Troubleshooting web interface issues](#) on page 83

Related tasks

[Creating and running a transition project](#) on page 69

Running prechecks

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

Before you begin

- You must have created a transition project.

About this task

You can run prechecks any number of times from Dashboard.

Steps

1. From **Dashboard**, select the subproject on which you want to run the prechecks.

The precheck message and the Precheck Status pane displays the results of the previous precheck operation.

2. Click **Run Prechecks**.

After the prechecks are complete, the results summary is displayed in the Precheck Status pane.

Note: The prechecks might take several minutes to complete.

3. Click **See Precheck Results** to view the results.

4. Choose an option under **Apply Type Filter** to filter the issues.

5. Choose **Security Only** under **Apply Tag Filter** to filter only security related issues.

- 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 display all messages related to security, select **Error**, **Warning**, **Informational**, and **Security Only**

Note: To view the results later, you must save the HTML file. You can also save the raw results in the CSV (comma separated values) format by clicking on the **csv file** link.

6. To export the results, click **Save as CSV**.

After you finish

You must resolve all the errors, and review the notifications that are detected by the prechecks before you start a transition.

Related concepts

[Severity levels for precheck messages](#) on page 49

[Requirements and restrictions for using the web interface](#) on page 66

Related tasks

[Creating a transition project](#) on page 72

[Starting a subproject](#) on page 75

[Completing a subproject](#) on page 76

Editing a subproject

You can edit a subproject to add or remove the volumes that are included, and to modify the transition configuration.

Before you begin

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

About this task

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

Steps

1. From **Dashboard**, select the subproject that you want to edit.
2. Click **Edit Configuration**.

The Volume selection pane of the New Transition Project wizard is displayed.

Note: Using the web interface you can remove the volumes only during the preparation phase.

3. Complete the wizard to edit your subproject.

Related concepts

[Requirements and restrictions for using the web interface](#) on page 66

[Restrictions on creating and configuring projects](#) on page 66

Related tasks

[Creating a transition project](#) on page 72

Managing logical interfaces

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

Before you begin

- You must have created a project.

Steps

1. Select a subproject from **Dashboard**.
2. To edit an existing logical interface, perform the following steps:
 - a. Click the **Edit** icon in the Network Address to be Migrated table.
The Modify LIF dialog box is displayed.
 - b. Make the required changes and click **Save**.
3. To remove an existing logical interface, click the remove icon in the Network Address to be Migrated table.

Note: You can remove a logical interface only when the subproject is in preparation, copy-baseline, copy-pause, or copy-update state. However, you cannot remove a logical interface if the cutover operation has failed.

4. To add a new logical interface, perform the following steps:
 - a. Click **Add LIF** below the Network Address to be Migrated table.
The **Add LIF** dialog box is displayed.
 - b. Select either **Add New LIF** (to add a new network interface) or **Transition existing 7-mode IP** (to use an existing network interface).
 - c. Enter the required values.
 - d. Click **Save** to add the new interface.

Related concepts

[Guidelines for deciding when to perform cutover](#) on page 47

Removing a volume or LIF from a subproject

If your subproject fails due to volume or LIF failure, you can remove the volumes and LIFs from the subproject during the transition process.

Generally, it is not recommended to use the tool CLI to perform operations on the subprojects created using the tool web interface.

You should use the CLI interface for volume removal operation if you have crossed the preparation phase. You should use the CLI interface for LIF removal operation if you are in the cutover phase. However the removal of the volume or LIF will be reflected on the Dashboard of the web interface only after clicking on a subproject to refresh the subproject contents.

Pausing and resuming a subproject

You can pause and resume a subproject that has started.

Before you begin

- You must have started a subproject.

About this task

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

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

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

Steps

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

Related concepts

[Requirements and restrictions for using the web interface](#) on page 66

Related tasks

[Creating a transition project](#) on page 72

[Starting a subproject](#) on page 75

[Aborting a subproject](#) on page 81

[Completing a subproject](#) on page 76

Aborting a subproject

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

Before you begin

The baseline copy must have started for the subproject.

About this task

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

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

Steps

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

2. Click **Abort**.

Related concepts

[Requirements and restrictions for using the web interface](#) on page 66

Related tasks

[Creating a transition project](#) on page 72

[Starting a subproject](#) on page 75

[Completing a subproject](#) on page 76

[Pausing and resuming a subproject](#) on page 80

[Deleting a project](#) on page 82

Deleting a project

You can delete a project from Dashboard.

Before you begin

- You must have created a project.

Steps

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

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

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

Related tasks

[Aborting a subproject](#) on page 81

Collecting tool logs

The 7-Mode Transition Tool creates log files that provide processing details of the transition operations that have occurred on your system. You can collect, download, and send tool logs by using the **Logs** button on the Home page.

Steps

1. Click **Logs** on the **Home** page
2. On the displayed **Transition logs** page, click **Collect tool logs** to collect the logs.

3. In the displayed **Transition Logs** dialog box, click **Yes** to download the logs to your local system.

The logs are downloaded in ZIP format, and the folder name is the timestamp.

Related concepts

[Troubleshooting web interface issues](#) on page 83

Troubleshooting web interface issues

You need to be aware of some of the common issues with the web interface and the steps to resolve them.

You cannot add credentials to controllers or clusters during transition

The tool displays the `Invalid credentials or cluster/controller already added.` error message when you try to add a controller or cluster with a password that has been changed.

Workaround

Using the tool web interface, remove the controller or cluster and add them again using the new credentials.

If the controller or cluster is part of an active project, you must restart the tool. The tool will then prompt you for information about controllers and clusters that are part of active projects.

Cannot select a secondary volume from the Volume selection pane in some scenarios

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

Workaround

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

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

Cannot select a volume with a broken SnapMirror relationship for transition

This issue usually occurs if the volume you selected for transition with the SnapMirror relationship has a broken SnapMirror relationship.

Workaround

Fix the SnapMirror relationship and include it in the project.

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

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

Workaround

Select the volume after some time.

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

The tool displays the Vserver does not have aggregate assigned. Please assign some aggregates to the vservers and click on the refresh button error message on the Volume mapping pane. This issue occurs if the selected Vserver has no aggregate other than the node root aggregate.

Workaround

Add an aggregate to the Vserver in clustered Data ONTAP.

Windows Program Compatibility Assistant notification displayed while installing the tool

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

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

Transitioning volumes using the 7-Mode Transition Tool (CLI)

The 7-Mode Transition Tool CLI provides a CLI wizard and commands that enable you to transition 7-Mode volumes. You must perform some tasks manually in the CLI, such as selecting the volumes in a session and transitioning a volume SnapMirror relationship. The CLI provides certain additional capabilities, such as creating a new Vserver.

The tool interacts with the 7-Mode storage system to collect volume and IP address information, verifies the 7-Mode configuration, and interacts with the cluster to apply the configuration on the Vserver.

You must perform the following tasks manually:

- You must select the volumes for transition.
- For transitioning a volume SnapMirror relationship, you must ensure that the secondary volumes and the primary volumes are correctly grouped in different sessions.
You must ensure that the secondary volume transition is completed before the primary volume transition.

You can perform the following additional tasks:

- Creating a new Vserver
- Providing multiple paths for SnapMirror data copy
- Editing the LIF configuration after the session is started
- Transitioning Kerberos configurations


Scenarios for selecting volumes in a session

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

Bidirectional SnapMirror relationship

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

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



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

Volumes with CIFS configuration

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

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

SnapMirror relationship between a primary volume and multiple secondary volumes

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

Transitioning stand-alone volumes by using the transition wizard

You can create a stand-alone type of transition session, where you can add a 7-Mode volume or a set of 7-Mode volumes that are not in a SnapMirror relationship. Transitioning stand-alone volumes involves creating a session, verifying the configuration, copying data to the clustered Data ONTAP volumes, and performing a cutover.

Before you begin

- The 7-Mode system and cluster must be prepared for transition.
- You must have gathered the [required information](#) on page 100 before you start the transition wizard.
- The 7-Mode system and the cluster must be always reachable from the Windows system that hosts the 7-Mode Transition Tool during the transition process; otherwise transition fails. The 7-Mode Transition Tool uses the management IP addresses to communicate with the 7-Mode system and the cluster. If these IP addresses are in the down state at any point, transition fails.
- Local users and groups that are accessing CIFS shares must not exist. If you have volumes with local users and groups that are being used for CIFS access, transition of such volumes is not supported if the target cluster is 8.2.0 and earlier versions, and the session is stopped.

Steps

1. From the **Start** menu, click **All Programs > NetApp 7-Mode Transition Tool > NetApp 7-Mode Transition Tool (cli) 1.2** to open the 7-Mode Transition Tool.

The main menu is displayed.

2. Create a new session by entering **2**.

Example

```
Step: Main menu
=====
Select one of the following menu items:
1. Show all sessions
2. Create a new session
3. Select existing session
4. Exit
Enter your selection(1-4): 2
```

- a. Press Enter to select the type of session as standalone.

```
Select the type of the session {primary, secondary, standalone} [standalone]:
```

- b. Select the 7-Mode storage system and enter the administrative user name and password.

Example

```

Step: Select the 7-Mode storage system
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Enter the FQDN or IP address of the 7-Mode system: system1.example.com
Enter the administrative user-name for the 7-Mode system 'system1.example.com' [root]:
Enter the password for 'root@system1.example.com':
INFO: Login to 'system1.example.com' as 'root' is successful.

```

- c. Optional: If you have enabled the MultiStore license, select the vFiler unit for transition.

Example

```

....
Found 5 vFiler units on the 7-Mode system 'system1':
vFiler Name          vFiler Status
-----
vfiler0              running
testVF1              running
testVF2              running
testVF3              running
testVF4              running
The vFiler unit contains volumes and protocol configurations.
You can select volumes for transition to Cluster-Mode.

Select the vFiler unit for transition [vfiler0]:

```

- d. Optional: Select the multipath IP address for transition.

Example

```

Enter another IP address for setting up multipath to the 7-Mode system(Press 'Enter'
to skip):

```

3. Select the 7-Mode volumes for transition.

Example

```

Step: Select 7-Mode volumes
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Retrieving volume information from 7-Mode system 'vfiler0@system1'.
This might take a few seconds.

Found 3 volumes that are candidates for transition on the 7-Mode system 'vfiler0@system1':
VolumeName      Size      Used(%)      Snapmirror
-----
vol0             851.48MB    31%          -
finance          50.00MB     0%           -
Purchase         50.00MB     0%           -

Enter the comma separated list of volumes for transition: finance, Purchase

You selected 2 volume(s):
VolumeName      Size      Used(%)      Snapmirror
-----
finance          50.00MB     0%           -
Purchase         50.00MB     0%           -

```

```
Confirm your selection {yes, no} [yes]:
```

4. Run a precheck operation on the 7-Mode system and resolve errors and warnings.

- a. Optional: Enter a name for the session.
- b. Run precheck on the 7-Mode system.

The precheck report lists the errors, warnings, and informational messages.

```
Step: Run precheck
=====
Provide a name for this session [Session_20130310_231131]: session_finance
Do you want to run precheck on the session {yes, no} [yes]:

Verify feasibility of transition - Check if all the features and functionality
included in this session is supported in Cluster-Mode.

Starting 7-Mode Precheck for this session 'session_finance'. This might take several
minutes.

[1/11 ] Session status checks                [      Ok      ]
[2/11 ] Validating 7-Mode system information [      Ok      ]
[3/11 ] Validating 7-Mode volumes information [      Ok      ]
[4/11 ] WAFL prechecks                      [      Ok      ]
[5/11 ] UNIX users and groups prechecks     [      Ok      ]
[6/11 ] SnapMirror prechecks                [ Notifications ]
[7/11 ] NFS prechecks                      [ Warnings      ]
[8/11 ] Networking prechecks               [ Notifications ]
[9/11 ] CIFS prechecks                    [ Warnings      ]
[10/11] Name services prechecks             [ Warnings      ]
[11/11] SAN prechecks                     [      Ok      ]

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

 0 Errors - No action required
15 Warnings - Need your attention!
15 Notifications - Please review
86 Informational messages
```

- c. Review the precheck report and resolve all errors to continue with the transition process.

Transition cannot be continued if even one error is present. You must also review the warnings and notifications to assess the impact of transition before continuing.

Example

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

15 Warnings:
-----
30141: Checking whether router failover is configured on the 7-Mode storage
system [      Warning      ]
      'vfiler0@system1'.
      > Router failover is configured on the 7-Mode storage system. However, router
failover is not
      > supported on the cluster.

40532: Verifying if netgroup entry comparisons will consider the domainnames in the
search directive [      Warning      ]
      from the '/etc/resolv.conf' file ('nis.netgroup.domain_search.enable' option) on 7-
Mode storage system.
      > Considering the domainnames in search directives from the '/etc/resolv.conf' file
while
      > doing netgroup entry comparison is not supported in clustered Data ONTAP.
```

```
40542: Verifying if netgroup entry comparisons will consider the legacy SUNOS
compatible nisdomainname [   Warning   ]

To continue, press Enter. To quit, press q and Enter:
```

5. Select the target cluster and enter the information as prompted by the wizard.

Example

```
Step: Select target cluster
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Enter the FQDN or management IP address of the cluster: cluster1.example.com
Enter the administrative user-name for the cluster 'cluster1.example.com' [admin]:
Enter the password for 'admin@cluster1.example.com':
INFO: Login to 'cluster1.example.com' as 'admin' is successful.
INFO: The name of the target cluster is 'cluster1'.
```

6. Press Enter to confirm that there are no local users or groups that are used for CIFS access.

This dialog box is displayed only if the target cluster is 8.2.0. If you enter **yes**, the session is stopped because transition of local users and groups is not supported for 8.2.0 and earlier releases of clustered Data ONTAP.

Example

```
Transition of local users and groups is not supported. Do you have local users and groups
accessing CIFS shares or files on the transitioning volumes in the 7-Mode storage system
'system1.example.com'? {yes, no} [no]:
```

The session is created and saved with a temporary name. You can change the session name later.

7. Configure the target Vserver.
 - a. Select the target Vserver or provide information about the new Vserver that is to be created during the transition start operation.

Note: The Vserver that is created during the start operation does not have any configurations applied to it. The Vserver is fully configured and ready to use after the transition operation finishes.

Example

```
Step: Select target Vserver
=====
You can type 'back', 'exit', 'help' or '?' at any time.

In this step, you can select a Vserver from the list of existing Vservers in the
cluster or provide information about a
new Vserver, which will be created by the 'transition start' operation.

Retrieving Vserver information from the cluster 'cluster1'.
This might take a few seconds.
```

Vserver	Root volume	State
-----	-----	-----

```

vs3          rootvs3      running
vs2          vs2root      running

Select an existing 'data' Vserver or enter a new Vserver name that must be created
during 'transition start' operation: vs_finance
INFO: You selected a new Vserver name 'vs_finance'.

```

- b. If you have selected a new Vserver, select the aggregate on which the root volume of the Vserver must be created.

Example

```

Retrieving aggregate information from the cluster 'cluster1'.
This might take a few seconds.

Found 1 online aggregate(s) on this cluster:
Aggregate name      Size      Used(%)      Node
-----
aggr1               1008.45MB    6%           cluster1-01

Select Vserver ('vs_finance') root volume aggregate [aggr1]:
INFO: Successfully added the cluster and Vserver information to the session.

```

8. Configure the clustered Data ONTAP volumes.
 - a. Select the aggregates where the new clustered Data ONTAP volumes should be created.

Example

```

Step: Configure target volumes
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Retrieving aggregate information from the Cluster 'cluster1'.
This might take a few seconds.

Found 1 online aggregate(s) on this cluster:
Aggregate name      Size      Used(%)      Node
-----
aggr1               1008.45MB    6%           cluster-01

Enter the comma separated list of aggregates to be used for Cluster-Mode volume
creation [aggr1]:
INFO: Aggregates 'aggr1' will be used for volume placement.

7-Mode volume      Cluster-Mode volume Size      Aggregate  In-Session?  Status
-----
finance            finance        20.00MB    aggr1        no           To be created
Purchase            Purchase        20.00MB    aggr1        no           To be created

```

- b. Optional: From the menu that appears, modify the volumes added to the session or the aggregates hosting the volumes.
- c. Accept the settings and save the session information by entering **1**.

```

Select one of the following menu items:
1. Accept the settings and commit the information to the session.
2. Include, exclude, or modify volumes.
3. Select aggregates.
4. Skip and continue.
Enter your selection(1-4): 1

INFO: Successfully modified volume-pair <finance, finance> in the session.
INFO: Successfully modified volume-pair <Purchase, Purchase> in the session.

```

9. Configure the network information to be transitioned.

- a. Specify the 7-Mode IP addresses that you want to transition to the target Vserver.

Example

```
Step: Configure network information
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Do you want to transition 7-Mode IP addresses {yes, no} [yes]:

NOTICE : All transitioning IP addresses will be configured on the target Vserver after
they have been removed from the 7-Mode storage system as part of the transition
complete operation.

IP addresses on the 7-Mode storage system 'system1':
IP address      Netmask      Default Gateway  Port
-----
192.0.2.69      255.255.255.192  192.40.0.1      e0c
192.0.2.70      255.255.255.192  192.40.0.1      e0d

Enter the comma separated list of IPv4 addresses from the previous list(Press 'Enter'
to skip):192.0.2.69
```

- b. Specify any new IP addresses that you want to configure on the Vserver after transition.

You can also add IP addresses that exist on the Vserver to the session.

Example

```
Specify the comma (,) separated list of new IPv4 addresses which do not exist on the 7-
Mode storage system :192.0.2.65
```

- c. Optional: From the menu that appears, modify the LIFs added to the session.

Example

```
Review the IP address configuration for transition based on the network connectivity
between the 7-Mode and Cluster-Mode systems.

IP address Netmask      Default Gateway  Home Node  Home Port  @7Mode/@Session/Status
-----
192.0.2.69 255.255.255.192 192.40.0.1  cluster1-01 e0c        Yes/No/To be created
```

- d. Accept the settings and save the session information by entering **1**.

```
Select one of the following menu items:
1. Accept the settings and commit the information to the session.
2. Include,exclude,or modify network interfaces.
3. Skip and continue.
Enter your selection(1-3): 1
```

10. Set the data copy schedule to be used for SnapMirror updates.

Example

```
Step: Get the data copy schedule parameters
=====
```

You can type 'back', 'exit', 'help' or '?' at any time.

Fetching cron schedule information from the cluster 'cluster1'.

Found '5' cron schedules:

Cron schedule name	Description
-----	-----
weekly	Sun@0:15
8hour	@2:15,10:15,18:15
5min	@:
00,:05,:10,:15,:20,:25,:30,:35,:40,:45,:50,:55	@:
hourly	@:05
daily	@0:10

Enter a cron schedule that exists in the cluster and which should be used for SnapMirror updates [daily]:

INFO: Specified cron schedule added to the session.

11. Run the precheck operation on the 7-Mode system and cluster and resolve errors and warnings.

- Run precheck on the 7-Mode system and cluster.

Example

Step: Run precheck

=====

Verify feasibility of transition - Check if all the features and functionality

included in this session is supported in Cluster-Mode.

Additionally, check for potential conflicts in configuration between the 7-Mode and Cluster-Mode systems.

Starting full Precheck for this session 'session_finance'. This might take several minutes.

[1/14]	Session status checks	[	Ok	]
[2/14]	Validating 7-Mode system information	[	Ok	]
[3/14]	Validating 7-Mode volumes information	[	Ok	]
[4/14]	Validating Cluster-Mode system information	[	Ok	]
[5/14]	Validating Cluster-Mode volume information	[	Ok	]
[6/14]	Validating LIF information	[	Notifications	]
[7/14]	WAFL prechecks	[	Ok	]
[8/14]	UNIX users and groups prechecks	[	Ok	]
[9/14]	SnapMirror prechecks	[	Notifications	]
[10/14]	NFS prechecks	[	Warnings	]
[11/14]	Networking prechecks	[	Notifications	]
[12/14]	CIFS prechecks	[	Warnings	]
[13/14]	Name services prechecks	[	Warnings	]
[14/14]	SAN prechecks	[	Ok	]

Operation summary:

The 'precheck' operation is completed with warnings and notifications.

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

- Review the results of the precheck and resolve errors and warnings.
- ## 12. From the main menu, view the session status by entering 4.

Example

Step: View session status

=====

Session name : session_finance
 Session status : Preparation

```

                Session type : standalone
                Apply configuration : Yes
                7-Mode system FQDN : system1.example.com
                7-Mode version : NetApp Release 8.1 7-Mode: Fri Nov 16 11:17:42 IST
2012
                vFiler unit : vfiler0
                Multipath IP address : -
                Cluster-Mode system FQDN : cluster1.example.com
                Cluster-Mode system Version : NetApp Release clustered Data ONTAP 8.2.0: Mon Mar
04 23:37:16 PST 2013
                Vserver information : vs_finance:aggr1
                Volume-Pairs : finance:finance, Purchase:Purchase
                Vserver data LIFs : 192.0.2.69
                SnapMirror schedule : daily

7-Mode volume Cluster-Mode volume Aggregate Snapmirror state Cluster-Mode volume status
-----
finance         finance         aggr1      -/-      To be created
Purchase        Purchase        aggr1      -/-      To be created

To continue, press Enter:

```

The session moves to the preparation state.

13. Start the data copy for the volumes in the session.

- a. From the main menu, start the baseline transfer by entering **7**.

Example

```

Step: Start data copy
=====
Do you want to run the start (data copy) operation on this session? {yes, no} [yes]:

Running the start (data copy) operation on the session 'session_finance'.

[1/18 ] Session status checks                      [      Ok      ]
[2/18 ] Validating 7-Mode system information         [      Ok      ]
[3/18 ] Validating 7-Mode volumes information        [      Ok      ]
[4/18 ] Validating Cluster-Mode system information   [      Ok      ]
[5/18 ] Validating Cluster-Mode volume information   [      Ok      ]
[6/18 ] Validating LIF information                   [ Notifications ]
[7/18 ] WAFL prechecks                              [      Ok      ]
[8/18 ] UNIX users and groups prechecks              [      Ok      ]
[9/18 ] SnapMirror prechecks                        [ Notifications ]
[10/18] NFS prechecks                               [ Warnings      ]
[11/18] Networking prechecks                       [ Notifications ]
[12/18] CIFS prechecks                             [ Warnings      ]
[13/18] Name services prechecks                     [ Warnings      ]
[14/18] SAN prechecks                              [      Ok      ]
[15/18] Creating Cluster-Mode Vserver                [      Ok      ]
[16/18] Creating Cluster-Mode volumes                [      Ok      ]
[17/18] Establishing SnapMirror relationships between the 7-Mode and Cluster-Mode
volumes                                             [      Ok      ]
[18/18] Initializing SnapMirror relationships between the 7-Mode and Cluster-Mode
volumes                                             [      Ok      ]

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

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

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

```

The session is in the copy-baseline state.

- b. Verify the results of the start operation.
- c. Press Enter to view the session status.

After the baseline transfer is completed, the session state changes to copy update. SnapMirror updates are performed based on the schedule selected. This is the state where you prepare for a cutover.

Example

```

Next suggested steps:
-----
1. Use the 'transition job-results -j 0efc4c54-0860-48e2-88b6-05a0f62f4b59' command to
   see the results of this
   operation.
2. Keep monitoring the session status by selecting 'View session status' or by running
   the CLI 'transition job-results
   0efc4c54-0860-48e2-88b6-05a0f62f4b59'.
3. Once the session status is 'Copy-update', cut-over operation can be started by
   selecting 'Start cut-over' from the
   transition main menu or by running 'transition complete -s session_finance' command
   from the CLI.

Do you want to see the session status? {yes, no} [yes]:

      Session name : session_finance
      Session status : Copy-update
      Session type : standalone
      Apply configuration : Yes
      7-Mode system FQDN : system1.example.com
      7-Mode version : NetApp Release 8.1 7-Mode: Fri Nov 16 11:17:42
IST 2012
      vFiler unit : vfiler0
      Multipath IP address : -
      Cluster-Mode system FQDN : cluster1.example.com
      Cluster-Mode system Version : NetApp Release 8.2.0 Cluster-Mode: Mon Mar 04
23:37:16 PST 2013
      Vserver information : vs_finance:aggr1:Created
      Volume-Pairs : finance:finance, Purchase:Purchase
      Vserver data LIFs : 192.0.2.69
      SnapMirror schedule : daily

7-Mode volume Cluster-Mode volume Snapmirror state Comments
-----
finance        finance        snapmirrored/idle * Last transfer size: 944.00KB
                  * Last transfer duration: 12
seconds
                  * Lag time: 3 minutes 12 seconds

Purchase       Purchase       snapmirrored/idle * Last transfer size: 680.00KB
                  * Last transfer duration: 12
seconds
                  * Lag time: 3 minutes 12 seconds

SUMMARY:
-----
1. Baseline transfer from the 7-Mode volumes to the Cluster-Mode volumes has been
   completed successfully. The session
   status is 'Copy-update'.
2. Cluster-Mode volumes are being updated periodically with the 7-Mode volumes based
   on the data copy schedule provided.
3. You can start the 'complete(cut-over)' operation by using the transition CLI
   command 'transition complete -s
   session_finance' or by selecting 'Start cut-over' from the 'transition menu'.

To continue, press Enter:

```

14. Monitor the data copy status and disconnect the client access when you are ready to perform cutover.
15. Optional: If you have NFS and Kerberos configured on the 7-Mode system, perform the following steps on the Vserver before cutover:
 - a. Log in as administrator to the cluster.
 - b. Use the `system license add` command to add the NFS license.

Example

```
cluster1::> system license add -license-code BBBBBBBBBBBBBB
```

- c. Use the `vserver services kerberos-realm show` command to verify whether the clock skew between the Kerberos KDC and the Vserver is less than or equal to 5 minutes.

Example

```
cluster1::> vserver services kerberos-realm show -fields clock-skew
configname          clock-skew
-----
krb_realm_0         5
```

- d. If the clock skew is greater than 5 minutes, use the `vserver services kerberos-realm modify` command to modify the Kerberos realm configuration.

For more information, see the *Clustered Data ONTAP File Access Management Guide for NFS*.

- e. If you want to transition an Active Directory-based Kerberos server, use the `vserver services dns create` command to configure a DNS entry for the AD domain on the Vserver.

Example

```
cluster1::> vserver services dns create -vserver vs_finance -domains
nfs.lab.example.com -name-servers 192.0.2.143 -state enabled
```

For more information about setting up a DNS domain, see the *Clustered Data ONTAP Network Management Guide*.

- f. Use the `vserver services dns show` command to verify that a DNS entry exists for the AD domain on the Vserver.

Example

```
cluster1::> vserver services dns show -vserver vs_finance
Vserver: vs_finance
Domains: nfs.lab.example.com
Name Servers: 192.0.2.143
Enable/Disable DNS: enabled
Timeout (secs): 2
Maximum Attempts: 1
```

16. Optional: If you have CIFS configured on the 7-Mode system, perform the following manual tasks before cutover:

- a. Log in to the 7-Mode system.
- b. Depending on whether you want to preserve the CIFS server name or use a new name, perform one of the following actions on the 7-Mode system:

If you want to...	Then...
Preserve the CIFS server name	i. Disconnect all CIFS access to the 7-Mode system. ii. Use the <code>cifs terminate</code> command to terminate the CIFS service. iii. Use the <code>cifs setup</code> command to create the CIFS server with a different name. The new CIFS server must be created on the 7-Mode system for transitioning CIFS configurations; otherwise, transition of CIFS share-level ACLs fails. The CIFS server name must be unique in the domain. iv. Ensure that the CIFS server is up and running.
Use a new CIFS server name	Ensure that the CIFS server is up and running.

Attention: If you want to preserve the CIFS server name of the 7-Mode system on the Vserver, the number of volumes in the 7-Mode system or vFiler unit to be transitioned must be less than or equal to 20 so that all the 7-Mode volumes can be transitioned in a single subproject or session.

- c. Log in as administrator to the cluster.
- d. Use the `system license add` command to add the CIFS license.

Example

```
cluster1::> system license add -license-code AAAAAAAAAAAAAA
```

- e. Verify that at least one data LIF exists on the Vserver.

If a data LIF does not exist, you can use the `network interface create` command to create the data LIF. For more information about creating data LIFs, see the *Clustered Data ONTAP Network Management Guide*.

- f. Use the `vserver services dns create` command to configure DNS on the Vserver.

Note: If you want to set up the CIFS server on the Vserver by preserving the identity, you must add the same DNS server as that of the 7-Mode system.

Example

```
cluster1:> vserver services dns create -vserver vs_finance -domains
cifs.lab.example.com -name-servers 192.0.2.145 -state enabled
```

For more information about setting up a DNS domain, see the *Clustered Data ONTAP Network Management Guide*.

- g. Use the `vserver cifs create` command to create a CIFS server on the Vserver.

If you want to preserve the 7-Mode CIFS server name on the Vserver, you must use the original CIFS server name that is used on the 7-Mode system.

For more information about configuring CIFS server, see the *Clustered Data ONTAP File Access Management Guide for CIFS*.

Example

```
cluster1:> vserver cifs create -vserver vs_finance -name CIFS_SERVER1 -domain
cifs.lab.example.com
```

17. Start the cutover operation.

- a. Enter **yes** to start the cutover operation.

Example

```
Step: Start the cutover operation
=====
Do you want to run the complete (cut-over) operation on this session? {yes, no} [yes]:
```

- b. Optional: If you have NFS Kerberos configuration on the 7-Mode system that you want to transition, enter the details of the Kerberos server.

Example

The following example shows a 7-Mode system that has a UNIX-based Kerberos server:

```
NFS Unix Kerberos configured on 7-Mode filer, do you want to transition {yes, no}
[yes]: yes
Enter the FQDN or IP Address of the Key Distribution Center(KDC) Server: 192.0.2.170
Enter the administrative user-name for the Key Distribution Center(KDC) Server
'192.0.2.170' [root]:
Enter the password for 'root@192.0.2.170':
```

The following example shows a 7-Mode system that has an Active Directory-based Kerberos server:

```
NFS Microsoft Kerberos configured on 7-Mode filer, do you want to transition {yes, no}
[yes]: yes
Enter the FQDN of the Active Directory Server: AUTH.SEC.EXAMPLE.COM
Enter the IP address of the Active Directory Server 'AUTH.SEC.EXAMPLE.COM'
[192.0.2.171]:
Enter the administrative user-name for the Active Directory Server
```

```
'AUTH.SEC.EXAMPLE.COM' [administrator]:
Enter the password for 'administrator@AUTH.SEC.EXAMPLE.COM':
```

Note: The transition of Organizational Unit, in which the machine account of Microsoft Active Directory Kerberos needs to be created, to clustered Data ONTAP 8.2.1 is not supported. As part of transition, the AD Kerberos machine account is created as part of the default Organizational Unit in the AD server. You should manually modify the Kerberos configuration for each of the data LIFs transitioned as part of the session to specify the Organizational Unit in which the Microsoft Active Directory Kerberos machine account needs to be created. For more information on how to manage the Kerberos configuration for LIFs, see the *Clustered Data ONTAP File Access Management Guide for NFS*.

- c. Select the mode in which to run the complete operation.

The default mode for completing the transition is normal.

Example

This complete operation can be run in either of the following two modes.

Test mode : If you select test mode, all the 7-Mode IP addresses and volumes added to the session are not modified in any way. All the client applications can continue to access data from the 7-Mode storage system after the complete operation. All the LIFs added to the session will be created on the Vserver in the 'down' state.

Normal mode : If you select normal mode, all the 7-Mode IP addresses added to the session will be unconfigured on the 7-Mode storage system and created on the Vserver in the 'up' state. All the 7-Mode volumes added to the session are made offline.

Select the mode in which you want to run this complete operation {normal, test}
[normal]:

- d. Verify the results of the complete operation and rectify any warnings.

18. From the main menu, verify that the session state is complete by entering 4.

Example

```
Step: Main menu
=====
Current session : session_finance
Select one of the following menu items:
1. Show all sessions
2. Create a new session
3. Select existing session
4. View session status
5. Exit
Enter your selection(1-5): 4

Step: View session status
=====
                Session name : session_finance
                Session status : Complete
                Session type  : standalone
Apply configuration : Yes
                7-Mode system FQDN : system1.example.com
                7-Mode version  : NetApp Release 8.1 7-Mode: Tue Oct 16 10:16:50 IST
```

```
2012
vFiler unit : vfiler0
Multipath IP address : -
Cluster-Mode system FQDN : cluster1.example.com
Cluster-Mode system Version : NetApp Release 8.2.0 Cluster-Mode: Tue Oct 16
14:16:59 IST 2012
Vserver information : vs_finance:aggr1:Created
Volume-Pairs : finance:finance, Purchase:Purchase
Vserver data LIFs : -
SnapMirror schedule : daily

7-Mode volume      Cluster-Mode volume  Snapmirror state      Comments
-----
finance            finance              -/-                   -
Purchase           Purchase             -/-                   -
```

After you finish

Perform any post-transition tasks before enabling client access to the Vserver.

Related concepts

- [Preparing for transition](#) on page 12
- [How you transition a stand-alone volume](#) on page 33
- [Severity levels for precheck messages](#) on page 49
- [Guidelines for deciding when to perform cutover](#) on page 47
- [Mode types you can use for completing the transition](#) on page 41
- [Prerequisites for transitioning NFS and CIFS configurations](#) on page 20
- [Error codes and corrective tasks for transition](#) on page 167

Related information

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

Transition setup worksheet for a stand-alone volume

Before you transition stand-alone volumes, you must gather the required information to complete the transition wizard for a stand-alone session.

You can use this worksheet to record the values that you require during the transition setup process. In the wizard, some fields might have a default value. You can either select the default value or enter your own value.

Session information

Information	Your value
Session name	

7-Mode information

Information	Your values
7-Mode storage system	
FQDN or IP address of the 7-Mode system	
User name	
Password	
vFiler unit, if any	
IP address for multipath	
7-Mode volumes	
List of volumes	
Network configuration information	
IP addresses to be transitioned	
New IP addresses to be configured on the Vserver	
Netmask	
Default gateway	
Route	

Cluster information

Information	Your values
Target cluster	
FQDN or cluster-management IP	
User name	
Password	
Target Vserver	
Vserver name	
Aggregate name For a new Vserver, the aggregate on which you want to create the Vserver's root volume	
Vserver audit logs destination path	

Information	Your values
Target volumes	
List of aggregates to be used for clustered Data ONTAP volume creation	
Data copy schedule parameters	
Data copy schedule on the cluster that is used to update the clustered Data ONTAP volumes	

NFS configuration information

If you have Kerberos authentication configured on the 7-Mode system, you need to enter the Kerberos server details during cutover when the 7-Mode Transition Tool prompts you. The tool detects whether the Kerberos Key Distribution Center (KDC) type is Active Directory-based or UNIX-based.

Information	Your values
Active Directory-based Kerberos KDC	
FQDN of the Active Directory server	
IP address of the Active Directory server	
User name	
Password	
UNIX-based Kerberos KDC	
IP address of the KDC server	
User name	
Password	

CIFS configuration information

If you have CIFS configured on the 7-Mode system, you need to manually configure DNS and CIFS server on the Vserver before the cutover operation.

Information	Your values
DNS domain	
Name server IP addresses	
CIFS server name	

Information	Your values
Active Directory domain	
Vserver name	

Transitioning a volume SnapMirror relationship using the transition wizard

To transition a 7-Mode volume SnapMirror relationship, you must transition the secondary volume before the primary volume. When the secondary transition is completed, a SnapMirror relationship is established between the 7-Mode primary and the clustered Data ONTAP secondary volumes. When the primary volume transition is completed, a SnapMirror relationship is established between the clustered Data ONTAP primary and secondary volumes.

Steps

1. [Transition setup worksheet for a SnapMirror relationship](#) on page 103
2. [Transitioning secondary volumes by using the transition wizard](#) on page 107
3. [Transitioning primary volumes by using the transition wizard](#) on page 119

Related concepts

[Preparing for transition](#) on page 12

[How you transition volumes in a SnapMirror relationship](#) on page 36

Transition setup worksheet for a SnapMirror relationship

Before you transition a volume SnapMirror relationship, you must gather the required information for completing the wizard for the primary and secondary session types.

You can use this worksheet to record the values that you require during the transition setup process. You can either accept the default value or enter your own value for the information types.

Session information

Information	Your value
Session name for secondary	
Session name for primary	

Secondary session

You must gather the following information about the 7-Mode environment:

Information	Your values
7-Mode storage system	
FQDN or IP address of the 7-Mode secondary system	
User name	
Password	
vFiler unit, if any	
IP address for multipathing	
7-Mode volumes	
List of volumes	
Network information	
IP addresses to be transitioned	
New IP addresses to be configured on the Vserver	
Netmask	
Default gateway	
SnapMirror primary 7-Mode storage system details	
FQDN or IP address of the 7-Mode primary system	
User name	
Password	
SnapMirror source data copy IP address	
Multipath IP address	

You must gather the following information about the cluster and Vserver:

Information	Your values
Target cluster	
FQDN or cluster-management IP	
User name	
Password	

Information	Your values
Target Vserver	
Vserver name	
Aggregate name For a new Vserver, the aggregate on which you want to create the Vserver's root volume	
Target volumes	
List of aggregates to be used for clustered Data ONTAP volume creation	
Data copy schedule parameters	
Data copy schedule on the cluster that is used to update the clustered Data ONTAP secondary volumes	

Primary session

You must gather the following information about the 7-Mode environment:

Information	Your values
7-Mode storage system	
FQDN or IP address of the 7-Mode primary system	
User name	
Password	
vFiler unit, if any	
IP address for multipathing	
7-Mode volumes	
List of volumes	
Network information	
IP addresses to be transitioned	
New IP addresses to be configured on the Vserver	
SnapMirror secondary cluster details	

Information	Your values
FQDN or IP address of the clustered Data ONTAP secondary system	
Vserver name	
User name	
Password	

You must gather the following information about the cluster and Vserver:

Information	Your values
Target cluster	
FQDN or cluster-management IP	
User name	
Password	
Target Vserver	
Vserver name	
Aggregate name For a new Vserver, the aggregate on which you want to create the Vserver's root volume	
Target volumes	
List of aggregates to be used for creating the clustered Data ONTAP volume	
Data copy schedule parameters	
Data copy schedule on the cluster that is used to update the clustered Data ONTAP primary volumes	

NFS configuration information

If you have Kerberos authentication configured on the 7-Mode system, you have to enter the Kerberos server details during cutover when the 7-Mode Transition Tool prompts you. The tool detects whether the Kerberos Key Distribution Center (KDC) type is Active Directory based or UNIX-based.

Information	Your values
Active Directory based Kerberos KDC	
FQDN of the Active Directory server	
IP address of the Active Directory server	
User name	
Password	
UNIX-based Kerberos KDC	
IP address of the KDC server	
User name	
Password	

CIFS configuration information

If you have CIFS configured on the 7-Mode system, you have to manually configure DNS and CIFS server on the Vserver before the cutover operation.

Information	Your values
DNS domain	
Name server IP addresses	
CIFS server name	
Active Directory domain	
Vserver name	

Transitioning secondary volumes by using the transition wizard

Transitioning a secondary volume involves creating a session, verifying the configuration, copying data to the clustered Data ONTAP volumes, and setting up a SnapMirror relationship between the 7-Mode primary volume and the clustered Data ONTAP secondary volume.

Before you begin

- The 7-Mode system and cluster must be prepared for transition.
- You must have gathered the *required information* before you start the transition wizard.
- The 7-Mode systems and clusters that are part of the transition must be always reachable from the 7-Mode Transition Tool during the transition process.

The 7-Mode Transition Tool uses the management IP addresses to reach the 7-Mode system and the cluster. If the IP addresses are in the down state at any point, transition fails.

- Local users and groups that are accessing CIFS shares must not exist.
If you have volumes with local users and groups that are being used for CIFS access, transition of such volumes is not supported if the target cluster is 8.2.0 and earlier versions, and the session is stopped.

Steps

- From the **Start** menu, click **All Programs > NetApp 7-Mode Transition Tool > NetApp 7-Mode Transition Tool (cli) 1.2** to open the 7-Mode Transition Tool.

The transition main menu is displayed.

- Create a new session by entering **2**.

Example

```
Step: Main menu
=====
Select one of the following menu items:
1. Show all sessions
2. Create a new session
3. Select existing session
4. Exit
Enter your selection(1-4): 2
```

- Enter **secondary** to select the session type .

Example

```
Select the type of the session {primary, secondary, standalone} [standalone]:secondary
```

- Select the 7-Mode storage system and enter the administrative user name and password.

Example

```
Step: Select the 7-Mode storage system
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Enter the FQDN or IP address of the secondary 7-Mode system: host4.example.com
Enter the administrative user-name for the secondary 7-Mode system 'host4.example.com':
[root]:
Enter the password for 'root@host4.example.com':
INFO: Login to 'host4.example.com' as 'root' is successful.
```

- Optional: If you have enabled the MultiStore license, select the vFiler unit for transition.

Example

```
....
Found 5 vFiler units on the 7-Mode system 'host4.example.com':
vFiler Name                vFiler Status
-----
vfiler0                    running
testVF1                   running
```

```
testVF2                      running
testVF3                      running
testVF4                      running
The vFile unit contains volumes and protocol configurations.
You can select volumes for transition to Cluster-Mode.
```

```
Select the vFile unit for transition [vfiler0]:
```

- d. Optional: Select the multipath IP address for transition.

Example

```
Enter another IP address for setting up multipath to the 7-Mode system(Press 'Enter'
to skip):
```

- e. Select the 7-Mode volumes for transition.

Example

```
Step: Select 7-Mode volumes
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Retrieving volume information from 7-Mode system 'host4.example.com'.
This might take a few seconds.

Found 5 volumes that are candidates for transition on the 7-Mode system 'host4':
VolumeName      Size      Used(%)      Snapmirror
-----
vol0             851.48MB    33%          -
Purchase_bk      20.00MB    1%           secondary
testVol1         20.00MB    1%           primary
Sales_bk         20.00MB    1%           secondary,primary
testVol          20.00MB    1%           primary

Enter the comma separated list of volumes for transition: Purchase_bk

You selected 1 volume(s):
VolumeName      Size      Used(%)      Snapmirror
-----
Purchase_bk      20.00MB    1%           secondary

Confirm your selection {yes, no} [yes]:
```

- f. Run a precheck operation on the 7-Mode system.

The precheck report lists errors, warnings, and informational messages.

Example

```
Step: Run precheck
=====
Provide a name for this session [Session_20121213_143318]: session_purchase
Do you want to run precheck on the session {yes, no} [yes]:

Verify feasibility of transition - Check if all the features and functionality
included in
this session is supported in Cluster-Mode.

Starting 7-Mode Precheck for this session 'session_purchase'. This might take several
minutes.

[1/11 ] Session status checks                      [      Ok      ]
[2/11 ] Validating 7-Mode system information        [      Ok      ]
```

```

[3/11 ] Validating 7-Mode volumes information      [      Ok      ]
[4/11 ] WAFL prechecks                            [      Ok      ]
[5/11 ] UNIX users and groups prechecks           [      Ok      ]
[6/11 ] SnapMirror prechecks                      [ Notifications ]
[7/11 ] NFS prechecks                            [   Warnings   ]
[8/11 ] Networking prechecks                    [ Notifications ]
[9/11 ] CIFS prechecks                          [      Ok      ]
[10/11] Name services prechecks                  [   Warnings   ]
[11/11] SAN prechecks                           [      Ok      ]

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

  0 Errors - No action required
 10 Warnings - Need your attention!
 13 Notifications - Please review
 68 Informational messages

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

```

- g. Verify the precheck report and resolve all the errors to continue with the transition process.
Transition cannot continue if even one error is present.
- h. Select the secondary target cluster and enter the information as prompted by the wizard.

Example

```

Step: Select target cluster
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Enter the FQDN or management IP address of the secondary cluster: cluster3.example.com
Enter the administrative user-name for the secondary cluster 'cluster3.example.com'
[admin]:
Enter the password for 'admin@cluster3.example.com':
INFO: Login to 'cluster3.example.com' as 'admin' is successful.
INFO: The name of the target cluster is 'cluster3'.
INFO: Using '192.0.2.115' as the cluster-management LIF.

```

- i. Press Enter to confirm there are no local users or groups.

If you enter **yes**, the session is stopped because transition of local users and groups is not supported for 8.2.0 and earlier releases of clustered Data ONTAP.

Example

```

Transition of local users and groups is not supported. Do you have local users and
groups
accessing CIFS shares or files on the transitioning volumes in the 7-Mode storage
system 'system1.example.com'? {yes, no} [no]:

```

- j. Select the target Vserver or provide information about the new Vserver to be created during the transition start operation.

Example

```

Step: Select target Vserver
=====
You can type 'back', 'exit', 'help' or '?' at any time.

In this step, you can select a Vserver from the list of existing Vservers in the

```

cluster or provide information about a new Vserver, which will be created by the 'transition start' operation.

Retrieving Vserver information from the cluster 'cluster3'.
This might take a few seconds.

Vserver	Root volume	State
-----	-----	-----
vssec	rootVolume	running

Select an existing 'data' Vserver or enter a new Vserver name that must be created during 'transition start' operation: vssec

INFO: You selected an existing Vserver 'vssec'.

INFO: Successfully added the cluster and Vserver information to the session.

- k. Configure the target volumes by selecting the aggregates where the new clustered Data ONTAP volumes should be created.

Note: You must ensure that there is enough space in the aggregates to create the volumes.

Example

Step: Configure target volumes

=====

You can type 'back', 'exit', 'help' or '?' at any time.

Retrieving aggregate information from the Cluster 'cluster3'.
This might take a few seconds.

Found 4 online aggregate(s) on this cluster:

Aggregate name	Size	Used(%)	Node
-----	----	-----	----
aggr3	224.12MB	0%	cluster3-01
aggr4	168.08MB	0%	cluster3-01
aggr1	56.04MB	73%	cluster3-01
aggr2	56.04MB	73%	cluster3-01

Enter the comma separated list of aggregates to be used for Cluster-Mode volume creation [aggr3]:

INFO: Aggregates 'aggr3' will be used for volume placement.

7-Mode volume	Cluster-Mode volume	Size	Aggregate	In-Session?	Status
-----	-----	----	-----	-----	-----
Purchase_bk	Purchase_bk	20.00MB	aggr3	no	To be created

Select one of the following menu items:

1. Accept the settings and commit the information to the session.
2. Include,exclude,or modify volumes.
3. Select aggregates.
4. Skip and continue.

Enter your selection(1-4): 1

INFO: Successfully modified volume-pair <Purchase_bk, Purchase_bk> in the session.

- l. From the menu that appears, accept the settings and save the session information by entering 1.

Example

Select one of the following menu items:

1. Accept the settings and commit the information to the session.
2. Include,exclude,or modify volumes.
3. Select aggregates.
4. Skip and continue.

```
Enter your selection(1-4): 1
INFO: Successfully modified volume-pair <Purchase_bk, Purchase_bk> in the session.
```

- m. Configure the network information by specifying the IP addresses that you want to transition to the target Vserver.

```
Step: Configure network information
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Do you want to transition 7-Mode IP addresses {yes, no} [yes]:

NOTICE : All transitioning IP addresses will be configured on the target Vserver after
they have been removed from the
7-Mode storage system as part of the transition complete operation.

IP addresses on the 7-Mode storage system 'host4':
IP address      Netmask      Default Gateway      Port
-----
192.0.2.75      255.255.255.192      192.40.0.1           e0c

Enter the comma separated list of IPv4 addresses from the previous list(Press 'Enter'
to skip):

Specify the comma separated list of new IPv4 addresses that do not exist on the 7-Mode
storage system(Press 'Enter' to skip):
192.0.2.67

Selected IP addresses:
1) 192.0.2.67

Confirm your selection {yes, no} [yes]:

Collecting LIF configurations of '192.0.2.67'
Enter the network mask: 255.255.255.192
Enter default gateway for the IP address '192.0.2.67': 192.40.0.1

Review the IP address configuration for transition based on the network connectivity
between the 7-Mode and Cluster-Mode systems.

IP address Netmask      Default Gateway Home Node      Home Port @7Mode/@Session/Status
-----
192.0.2.67 255.255.255.192  192.40.0.1      cluster3-01 e0c          No/No/To be created
```

- n. Add the details of the 7-Mode primary system.

Example

```
Step: Add details of the 7-Mode primary system
=====
You can type 'back', 'exit', 'help' or '?' at any time.

List of SnapMirror relationships between 7-Mode primary and secondary, which can be
added to the session are:

Source Host      Source Volumes      Destination Host      Destination Volumes
-----
192.0.2.123      Purchase             host4.example.com      Purchase_bk

Do you want to add details for the 7-Mode primary system '192.0.2.123' {yes, no} [yes]:
Enter the FQDN or IP address of the 7-Mode primary system (or the default vFiler unit
if Multistore is licensed) [192.0.2.123]: host5.example.com
Enter the administrative user-name for the 7-Mode system 'host5.example.com' [root]:
Enter the password for 'root@host5.example.com':
INFO: Login to 'host5.example.com' as 'root' is successful.
Enter the SnapMirror source data copy IPv4 address [192.0.2.123]:
```

- o. Set the data copy schedule to be used for SnapMirror updates.

Example

```
Step: Get the data copy schedule parameters
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Fetching cron schedule information from the cluster 'cluster3'.
Found '7' cron schedules:
Cron schedule name          Description
-----
weekly                      Sun@0:15
8hour                      @2:15,10:15,18:15
nightly                     @0:00
4hour                      @8:00,12:00,16:00,20:00
5min                       @:
00,:05,:10,:15,:20,:25,:30,:35,:40,:45,:50,:55
daily                      @0:10
avUpdateSchedule           @2:00

Enter a cron schedule that exists in the cluster and which should be used for
SnapMirror updates [daily]:
INFO: Specified cron schedule added to the session.
```

- p. View the results of the precheck on the 7-Mode system and cluster and resolve errors, if any.

The transition session is in the preparation state.

3. Start the data copy by entering 7.

Example

```
Step: Start data copy
=====
Do you want to run the start (data copy) operation on this session? {yes, no} [yes]:

Running the start (data copy) operation on the session 'session_purchase'.

[1/17 ] Session status checks                      [      Ok      ]
[2/17 ] Validating 7-Mode system information        [      Ok      ]
[3/17 ] Validating 7-Mode volumes information       [      Ok      ]
[4/17 ] Validating Cluster-Mode system information [      Ok      ]
[5/17 ] Validating Cluster-Mode volume information [      Ok      ]
[6/17 ] Validating LIF information                  [      Ok      ]
[7/17 ] WAFL prechecks                             [      Ok      ]
[8/17 ] UNIX users and groups prechecks             [      Ok      ]
[9/17 ] SnapMirror prechecks                       [ Warnings     ]
[10/17] NFS prechecks                              [ Warnings     ]
[11/17] Networking prechecks                       [ Notifications]
[12/17] CIFS prechecks                             [      Ok      ]
[13/17] Name services prechecks                    [ Warnings     ]
[14/17] SAN prechecks                              [      Ok      ]
[15/17] Creating Cluster-Mode volumes               [      Ok      ]
[16/17] Establishing SnapMirror relationships between the 7-Mode and Cluster-Mode volumes [      Ok      ]
[17/17] Initializing SnapMirror relationships between the 7-Mode and Cluster-Mode volumes [      Ok      ]

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

 0 Errors - No action required
11 Warnings - Need your attention!
```

```
13 Notifications - Please review
81 Informational messages
```

The tool performs another precheck during data copy. The session is in the copy-baseline state.

4. Optional: View the session status and prepare for cutover.

After the baseline transfer is completed, the session state changes to copy update. SnapMirror updates are performed based on the schedule selected. This is the state where you prepare for a cutover.

Example

```

                Session name : Session_20121213_163606
                Session status : Copy-update
                Session type : secondary
                Apply configuration : Yes
                7-Mode system FQDN : host4.example.com
                7-Mode version : NetApp Release 8.1 7-Mode: Sun Nov 18 11:15:50 IST
2012
                vFiler unit : vfiler0
                Multipath IP address : -
                Cluster-Mode system FQDN : 192.0.2.115
                Cluster-Mode system Version : NetApp Release 8.2.0 Cluster-Mode: Wed Dec 05
01:45:54 IST 2012
                Vserver information : vssec:Pre-existing
                Volume-Pairs : Purchase_bk:Purchase_bk
                Vserver data LIFs : 192.0.2.67
                SnapMirror schedule : daily
                Primary 7-Mode details : 192.0.2.123:host5.example.com:192.0.2.123:-

7-Mode volume Cluster-Mode volume Snapmirror state  Comments
-----
Purchase_bk   Purchase_bk           snapmirrored/idle  * Last transfer size: 400.00KB
                                           * Last transfer duration: 9 seconds
                                           * Lag time: 24 days 21 hours 0
minutes 16 seconds
```

5. Monitor the data copy status and disconnect the client access when you are ready to perform cutover.

6. Optional: If you have NFS and Kerberos configured on the 7-Mode system, perform the following steps on the Vserver before cutover:

- a. Log in as administrator to the cluster.
- b. Use the system license add command to add the NFS license.

Example

```
cluster1::> system license add -license-code BBBBBBBBBBBBBB
```

- c. Use the vsriver services kerberos-realm show command to verify whether the clock skew between the Kerberos KDC and the Vserver is less than or equal to 5 minutes.

Example

```
cluster1::> vsriver services kerberos-realm show -fields clock-skew
configname          clock-skew
```

```
-----
krb_realm_0      5
```

- d. If the clock skew is greater than 5 minutes, use the `vserver services kerberos-realm modify` command to modify the Kerberos realm configuration.

For more information, see the *Clustered Data ONTAP File Access Management Guide for NFS*.

- e. If you want to transition an Active Directory-based Kerberos server, use the `vserver services dns create` command to configure a DNS entry for the AD domain on the Vserver.

Example

```
cluster1::> vserver services dns create -vserver vs_finance -domains
nfs.lab.example.com -name-servers 192.0.2.143 -state enabled
```

For more information about setting up a DNS domain, see the *Clustered Data ONTAP Network Management Guide*.

- f. Use the `vserver services dns show` command to verify that a DNS entry exists for the AD domain on the Vserver.

Example

```
cluster1::> vserver services dns show -vserver vs_finance

      Vserver: vs_finance
      Domains: nfs.lab.example.com
      Name Servers: 192.0.2.143
      Enable/Disable DNS: enabled
      Timeout (secs): 2
      Maximum Attempts: 1
```

7. Optional: If you have CIFS configured on the 7-Mode system, perform the following manual tasks before cutover:
 - a. Log in to the 7-Mode system.
 - b. Depending on whether you want to preserve the CIFS server name or use a new name, perform one of the following actions on the 7-Mode system:

If you want to...	Then...
Preserve the CIFS server name	i. Disconnect all CIFS access to the 7-Mode system. ii. Use the <code>cifs terminate</code> command to terminate the CIFS service. iii. Use the <code>cifs setup</code> command to create the CIFS server with a different name. The new CIFS server must be created on the 7-Mode system for transitioning CIFS configurations; otherwise, transition of CIFS share-level ACLs fails. The CIFS server name must be unique in the domain. iv. Ensure that the CIFS server is up and running.
Use a new CIFS server name	Ensure that the CIFS server is up and running.

Attention: If you want to preserve the CIFS server name of the 7-Mode system on the Vserver, the number of volumes in the 7-Mode system or vFile unit to be transitioned must be less than or equal to 20 so that all the 7-Mode volumes can be transitioned in a single subproject or session.

- c. Log in as administrator to the cluster.
- d. Use the `system license add` command to add the CIFS license.

Example

```
cluster1::> system license add -license-code AAAAAAAAAAAAAA
```

- e. Verify that at least one data LIF exists on the Vserver.

If a data LIF does not exist, use the `network interface create` command to create the data LIF. For more information about creating data LIFs, see the *Clustered Data ONTAP Network Management Guide*.

- f. Use the `vserver services dns create` command to configure DNS on the Vserver.

Note: If you want to set up the CIFS server on the Vserver by preserving the identity, you must add the same DNS server as that of the 7-Mode system.

Example

```
cluster1::> vserver services dns create -vserver vs_finance -domains  
cifs.lab.example.com -name-servers 192.0.2.145 -state enabled
```

For more information about setting up a DNS domain, see the *Clustered Data ONTAP Network Management Guide*.

- g. Use the `vserver cifs create` command to create a CIFS server on the Vserver.

If you want to preserve the 7-Mode CIFS server name on the Vserver, you must use the original CIFS server name that is used on the 7-Mode system.

For more information about configuring CIFS server, see the *Clustered Data ONTAP File Access Management Guide for CIFS*.

Example

```
cluster1::> vsriver cifs create -vserver vs_finance -name CIFSSEVER1 -domain
cifs.lab.example.com
```

8. From the 7-Mode Transition Tool, confirm that the data copy is completed and start the cutover operation.
 - a. Enter **yes** to start the cutover operation.

Example

```
Step: Start the cutover operation
=====
Do you want to run the complete (cut-over) operation on this session? {yes, no} [yes]:
```

- b. Optional: If you have NFS Kerberos configuration on the 7-Mode system that you want to transition, enter the details of the Kerberos server.

Example

The following example shows a 7-Mode system that has a UNIX-based Kerberos server:

```
NFS Unix Kerberos configured on 7-Mode filer, do you want to transition {yes, no}
[yes]: yes
Enter the FQDN or IP Address of the Key Distribution Center(KDC) Server: 192.0.2.170
Enter the administrative user-name for the Key Distribution Center(KDC) Server
'192.0.2.170' [root]:
Enter the password for 'root@192.0.2.170':
```

The following example shows a 7-Mode system that has an Active Directory-based Kerberos server:

```
NFS Microsoft Kerberos configured on 7-Mode filer, do you want to transition {yes, no}
[yes]: yes
Enter the FQDN of the Active Directory Server: AUTH.SEC.EXAMPLE.COM
Enter the IP address of the Active Directory Server 'AUTH.SEC.EXAMPLE.COM'
[192.0.2.171]:
Enter the administrative user-name for the Active Directory Server
'AUTH.SEC.EXAMPLE.COM' [administrator]:
Enter the password for 'administrator@AUTH.SEC.EXAMPLE.COM':
```

Note: The transition of Organizational Unit, in which the machine account of Microsoft Active Directory Kerberos needs to be created, to clustered Data ONTAP 8.2.1 is not supported. As part of transition, the AD Kerberos machine account is created as part of the default Organizational Unit in the AD server. You should manually modify the Kerberos configuration for each of the data LIFs transitioned as part of the session to specify the Organizational Unit in which the Microsoft Active Directory Kerberos machine account

needs to be created. For more information on how to manage the Kerberos configuration for LIFs, see the *Clustered Data ONTAP File Access Management Guide for NFS*.

- c. Select the mode in which to run the complete operation.

The default mode for completing the transition is normal.

Example

This complete operation can be run in either of the following two modes.

Test mode : If you select test mode, all the 7-Mode IP addresses and volumes added to the session are not modified in any way. All the client applications can continue to access data from the 7-Mode storage system after the complete operation. All the LIFs added to the session will be created on the Vserver in the 'down' state.

Normal mode : If you select normal mode, all the 7-Mode IP addresses added to the session will be unconfigured on the 7-Mode storage system and created on the Vserver in the 'up' state. All the 7-Mode volumes added to the session are made offline.

Select the mode in which you want to run this complete operation {normal, test}
[normal]:

- d. Verify the results of the complete operation and rectify any errors or warnings.

9. Verify that the session state is complete by entering 4.

```
Step: Main menu
=====
Current session : session_finance
Select one of the following menu items:
1. Show all sessions
2. Create a new session
3. Select existing session
4. View session status
5. Exit
Enter your selection(1-5): 4
```

Example

The following example shows that the session state is Complete.

```
Step: View session status
=====
                Session name : Session_20121213_163606
                Session status : Complete
                Session type : secondary
                Apply configuration : Yes
                7-Mode system FQDN : host4.example.com
                7-Mode version : NetApp Release 8.1 7-Mode: Sun Nov 18 11:15:50 IST
2012
                vFiler unit : vfiler0
                Multipath IP address : -
                Cluster-Mode system FQDN : 192.0.2.115
                Cluster-Mode system Version : NetApp Release 8.2.0 Cluster-Mode: Wed Dec 05
01:45:54 IST 2012
                Vserver information : vssec:Pre-existing
                Volume-Pairs : Purchase_bk:Purchase_bk
                Vserver data LIFs : 192.0.2.67
```

SnapMirror schedule : daily			
Primary 7-Mode details : 192.0.2.123:host5.example.com:192.0.2.123:-			
7-Mode volume	Cluster-Mode volume	Snapmirror state	Comments
-----	-----	-----	-----
Purchase_bk	Purchase_bk	-/-	-

After you finish

Perform the post-transition tasks before enabling client access to the Vserver.

Related concepts

[Preparing for transition](#) on page 12

[How you transition volumes in a SnapMirror relationship](#) on page 36

[Severity levels for precheck messages](#) on page 49

[Guidelines for deciding when to perform cutover](#) on page 47

[Mode types you can use for completing the transition](#) on page 41

[Prerequisites for transitioning NFS and CIFS configurations](#) on page 20

[Error codes and corrective tasks for transition](#) on page 167

Related information

[Documentation on the NetApp Support Site: support.netapp.com](#)

Transitioning primary volumes by using the transition wizard

Transitioning a primary volume involves creating a session, verifying the configuration, copying data to the primary clustered Data ONTAP volumes, and setting up a SnapMirror relationship between the primary and secondary clustered Data ONTAP volumes.

Before you begin

- The 7-Mode system and cluster must be prepared for transition.
- You must have gathered the [required information](#) before you start the transition wizard.
- The 7-Mode system and the cluster must be always reachable from the 7-Mode Transition Tool during the transition process; otherwise transition fails.

The 7-Mode Transition Tool uses the management IP addresses to reach the 7-Mode system and the cluster. If these IP addresses are in the down state at any point, transition fails.

- Local users and groups that are accessing CIFS shares must not exist.
If you have volumes with local users and groups that are being used for CIFS access, transition of such volumes is not supported, if the target cluster is 8.2.0 and earlier versions, and the session is stopped.

Steps

1. From the **Start** menu, click **All Programs > NetApp 7-Mode Transition Tool > NetApp 7-Mode Transition Tool (cli) 1.2** to open the 7-Mode Transition Tool.

The transition main menu is displayed.

2. Create a new session by entering **2**.

Example

```
Step: Main menu
=====
Select one of the following menu items:
1. Show all sessions
2. Create a new session
3. Select existing session
4. Exit
Enter your selection(1-4): 2
```

- a. Enter **primary** to select the type of session as primary.

```
Select the type of the session {primary, secondary, standalone} [standalone]:primary
```

- b. Select the 7-Mode storage system and enter the administrative user name and password.

Example

```
Step: Select the 7-Mode storage system
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Enter the FQDN or IP address of the primary 7-Mode system: host1.example.com
Enter the administrative user-name for the primary 7-Mode system 'host1.example.com'
[root]:
Enter the password for 'root@host1.example.com':
INFO: Login to 'host1.example.com' as 'root' is successful.
```

- c. Optional: If you have enabled the MultiStore license, select the vFiler unit for transition.

Example

```
....
Found 5 vFiler units on the 7-Mode system 'host1.example.com':
vFiler Name                vFiler Status
-----
vfiler0                     running
testVF1                    running
testVF2                    running
testVF3                    running
testVF4                    running
The vFiler unit contains volumes and protocol configurations.
You can select volumes for transition to Cluster-Mode.

Select the vFiler unit for transition [vfiler0]:
```

- d. Optional: Select the multipath IP address for transition.

Example

```
Enter another IP address for setting up multipath to the 7-Mode system(Press 'Enter'
to skip):
```

- e. Select the 7-Mode volumes for transition.

Example

```

Step: Select 7-Mode volumes
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Retrieving volume information from 7-Mode system 'vfiler0@host1.example.com'.
This might take a few seconds.

Found 7 volumes that are candidates for transition on the 7-Mode system
'vfiler0@host4.example.com':
VolumeName      Size      Used(%)      Snapmirror
-----
Sales            20.00MB      1%          fan-out(4)
dataVol5         20.00MB      1%          -
vol0             851.48MB     54%         -
rootVol5         20.00MB      1%          -
testVol1         20.00MB      1%          -
Purchase         20.00MB      1%          fan-out(3)
testVol          20.00MB      1%          -

Enter the comma separated list of volumes for transition: Purchase

You selected 1 volume(s):
VolumeName      Size      Used(%)      Snapmirror
-----
Purchase         20.00MB      1%          fan-out(3)

```

- f. Run a precheck operation on the 7-Mode system and optionally provide a name for the session.

The precheck report lists errors, warnings, and informational messages.

- g. Verify the precheck report and resolve all the errors to continue with the transition process.

Transition cannot continue if even one error is present.

- h. Select the target cluster and enter the information as prompted by the wizard.

Example

```

Step: Select target cluster
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Enter the FQDN or management IP address of the primary cluster: cluster1.example.com
Enter the administrative user-name for the primary cluster 'cluster1.example.com'
[admin]:
Enter the password for 'admin@cluster1.example.com':
INFO: Login to 'cluster1.example.com' as 'admin' is successful.
INFO: The name of the target cluster is 'cluster1'.
INFO: Using '192.0.2.126' as the cluster-management LIF.

```

- i. Press Enter to confirm there are no local users or groups.

If you enter **yes**, the session is stopped because transition of local users and groups is not supported for 8.2.0 and earlier releases of clustered Data ONTAP.

Example

```

Transition of local users and groups is not supported. Do you have local users and
groups
accessing CIFS shares or files on the transitioning volumes in the 7-Mode storage
system 'system1.example.com'? {yes, no} [no]:

```

- j. Select the target Vserver or provide information about the new Vserver that is to be created during the transition start operation.

Example

```
Step: Select target Vserver
=====
You can type 'back', 'exit', 'help' or '?' at any time.

In this step, you can select a Vserver from the list of existing Vservers in the
cluster or provide information about a
new Vserver, which will be created by the 'transition start' operation.

Retrieving Vserver information from the cluster 'cluster1'.
This might take a few seconds.

Vserver          Root volume      State
-----
vs39             rootVolume      running

Select an existing 'data' Vserver or enter a new Vserver name that must be created
during 'transition start' operation: vs39
INFO: You selected an existing Vserver 'vs39'.

INFO: Successfully added the cluster and Vserver information to the session.
```

- k. Configure the target volumes by selecting the aggregates where the new clustered Data ONTAP volumes should be created.

Example

```
Step: Configure target volumes
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Retrieving aggregate information from the Cluster 'cluster1'.
This might take a few seconds.

Found 2 online aggregate(s) on this cluster:
Aggregate name   Size      Used(%)    Node
-----
aggr2            1.47GB    0%         cluster1-01
aggr1            1.47GB    1%         cluster1-01

Enter the comma separated list of aggregates to be used for Cluster-Mode volume
creation [aggr2]:
INFO: Aggregates 'aggr2' will be used for volume placement.

7-Mode volume   Cluster-Mode volume   Size      Aggregate In-Session? Status
-----
Purchase        Purchase              20.00MB   aggr2      no           To be created
```

- l. From the menu that appears, accept the settings and save the session information by entering **1**.

Example

```
Select one of the following menu items:
1. Accept the settings and commit the information to the session.
2. Include, exclude, or modify volumes.
3. Select aggregates.
4. Skip and continue.
Enter your selection(1-4): 1
```

```
INFO: Successfully modified volume-pair <Purchase, Purchase> in the session.
```

- m. Configure the network information by specifying the IP addresses that you want to transition to the target Vserver.

```
Step: Configure network information
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Do you want to transition 7-Mode IP addresses {yes, no} [yes]:

NOTICE : All transitioning IP addresses will be configured on the target Vserver after
they have been removed from the
7-Mode storage system as part of the transition complete operation.

IP addresses on the 7-Mode storage system 'vfiler0@user1-vsrm2':
IP address          Netmask          Default Gateway      Port
-----
192.0.2.65          255.255.255.192      192.40.0.1           e0c

Enter the comma separated list of IPv4 addresses from the previous list(Press 'Enter'
to skip):

Specify the comma separated list of new IPv4 addresses that do not exist on the 7-Mode
storage system(Press 'Enter' to skip):
192.0.2.66

Selected IP addresses:
1) 192.0.2.66

Confirm your selection {yes, no} [yes]:

Collecting LIF configurations of '192.0.2.66'
Enter the network mask: 255.255.255.192
Enter default gateway for the IP address '192.0.2.66': 192.40.0.1

Review the IP address configuration for transition based on the network connectivity
between the 7-Mode and Cluster-Mode systems.

IP address Netmask   Default Gateway   Home Node   Home Port   @7Mode/@Session/Status
-----
192.0.2.66 255.255.255.192 192.40.0.1 cluster1-01 e0c         No/No/To be created
```

- n. Specify the SnapMirror secondary cluster details.

Example

```
Step: Add details of the Cluster-Mode secondary system
=====
You can type 'back', 'exit', 'help' or '?' at any time.

List of SnapMirror relations between 7-Mode primary and C-Mode secondary(possibly),
which are not yet added to the
session:

7-Mode primary   Primary volumes   Destination Vserver   Destination volumes
-----
user1-vsrm2      Purchase          vssec                 Purchase_bk
user1-vsrm2      Purchase          vs783                 Purchase_bk
user1-vsrm2      Purchase          user1-vsrm4           Purchase_bk

INFO: Entries in the 'Destination Vserver' column may not represent the Vserver. If it
is not a Vserver no need to add it to
the session.

Do you want to add details for the Vserver 'vssec' (Enter 'no' if it is not a Vserver)
{yes, no} [yes]:
Enter the FQDN or management IP address of the secondary cluster: cluster2.example.com
```

```

Enter the administrative user-name for the secondary cluster 'cluster2.example.com'
[admin]:
Enter the password for 'admin@cluster2.example.com':
INFO: Login to 'cluster2.example.com' as 'admin' is successful.
INFO: The name of the target cluster is 'cluster2'.
INFO: Using '192.0.2.66' as the cluster-management LIF.
INFO: Information about the secondary cluster has been added to the session.

```

- o. Provide the data copy schedule to be used for SnapMirror updates.

Example

```

Step: Get the data copy schedule parameters
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Fetching cron schedule information from the cluster 'cluster1'.
Found '5' cron schedules:
Cron schedule name                                Description
-----
weekly                                              Sun@0:15
CronJob_vs3_voltest_1                             @8:00,12:00,16:00,20:00
CronJob_vs3_voltest_0                             @0:00
CronJob_vs2_fin_data_7                             @8:00,12:00,16:00,20:00
daily                                              @0:10
hourly                                             @:05

Enter a cron schedule that exists in the cluster and which should be used for
SnapMirror updates [daily]:

```

- p. View the results of the precheck on the 7-Mode system and cluster and resolve errors, if any.

The transition session is in the preparation state.

3. Start data copy by entering 7.

Example

```

Step: Start data copy
=====
Do you want to run the start (data copy) operation on this session? {yes, no} [yes]:

Running the start (data copy) operation on the session 'session_purchase2'.

[1/17 ] Session status checks                      [      Ok      ]
[2/17 ] Validating 7-Mode system information        [      Ok      ]
[3/17 ] Validating 7-Mode volumes information       [      Ok      ]
[4/17 ] Validating Cluster-Mode system information [      Ok      ]
[5/17 ] Validating Cluster-Mode volume information [      Ok      ]
[6/17 ] Validating LIF information                  [ Notifications ]
[7/17 ] WAFL prechecks                             [      Ok      ]
[8/17 ] UNIX users and groups prechecks             [  Warnings    ]
[9/17 ] SnapMirror prechecks                       [  Warnings    ]
[10/17] NFS prechecks                              [  Warnings    ]
[11/17] Networking prechecks                       [ Notifications ]
[12/17] CIFS prechecks                             [      Ok      ]
[13/17] Name services prechecks                    [  Warnings    ]
[14/17] SAN prechecks                              [      Ok      ]
[15/17] Creating Cluster-Mode volumes               [      Ok      ]
[16/17] Establishing SnapMirror relationships between the 7-Mode and Cluster-Mode volumes

[      Ok      ]
[17/17] Initializing SnapMirror relationships between the 7-Mode and Cluster-Mode volumes

[      Ok      ]

Operation summary:

```

```
-----
The 'start' operation is completed with warnings and notifications.
```

```
 0 Errors - No action required
16 Warnings - Need your attention!
14 Notifications - Please review
78 Informational messages
```

The tool performs another precheck during data copy. The session is in the copy-baseline state.

4. Optional: View the session status.

After the baseline transfer is completed, the session state changes to copy update. SnapMirror updates are performed based on the schedule selected. This is state where you prepare for cutover.

Example

```

                Session name : session_purchase2
                Session status : Copy-update
                Session type : primary
                Apply configuration : Yes
                7-Mode system FQDN : host1.example.com
                7-Mode version : NetApp Release 8.1 7-Mode: Fri Nov 16 11:17:42 IST
2012
                vFiler unit : vfiler0
                Multipath IP address : -
                Cluster-Mode system FQDN : 192.0.2.126
                Cluster-Mode system Version : NetApp Release 8.2.0 Cluster-Mode: Wed Dec 12
01:47:08 IST 2012
                Vserver information : vs39:Pre-existing
                Volume-Pairs : Purchase:Purchase
                Vserver data LIFs : -
                SnapMirror schedule : -
                Secondary cluster details : -

7-Mode volume Cluster-Mode volume Snapmirror state Comments
-----
Purchase      Purchase      snapmirrored/idle * Last transfer size: 932.00KB
                                     * Last transfer duration: 12 seconds
                                     * Lag time: 8 minutes 31 seconds
```

5. Monitor the data copy status and disconnect the client access when you are ready to perform cutover.
6. Optional: If you have NFS and Kerberos configured on the 7-Mode system, perform the following steps on the Vserver before cutover:
 - a. Log in as administrator to the cluster.
 - b. Use the `system license add` command to add the NFS license.

Example

```
cluster1:> system license add -license-code BBBBBBBBBBBBBB
```

- c. Use the `vserver services kerberos-realm show` command to verify whether the clock skew between the Kerberos KDC and the Vserver is less than or equal to 5 minutes.

Example

```
cluster1:> vserver services kerberos-realm show -fields clock-skew
configname      clock-skew
-----
krb_realm_0     5
```

- d. If the clock skew is greater than 5 minutes, use the `vserver services kerberos-realm modify` command to modify the Kerberos realm configuration.

For more information, see the *Clustered Data ONTAP File Access Management Guide for NFS*.

- e. If you want to transition an Active Directory-based Kerberos server, use the `vserver services dns create` command to configure a DNS entry for the AD domain on the Vserver.

Example

```
cluster1:> vserver services dns create -vserver vs_finance -domains
nfs.lab.example.com -name-servers 192.0.2.143 -state enabled
```

For more information about setting up a DNS domain, see the *Clustered Data ONTAP Network Management Guide*.

- f. Use the `vserver services dns show` command to verify that a DNS entry exists for the AD domain on the Vserver.

Example

```
cluster1:> vserver services dns show -vserver vs_finance

Vserver: vs_finance
Domains: nfs.lab.example.com
Name Servers: 192.0.2.143
Enable/Disable DNS: enabled
Timeout (secs): 2
Maximum Attempts: 1
```

7. Optional: If you have CIFS configured on the 7-Mode system, perform the following manual tasks before cutover:
 - a. Log in to the 7-Mode system.
 - b. Depending on whether you want to preserve the CIFS server name or use a new name, perform one of the following actions on the 7-Mode system:

If you want to...	Then...
Preserve the CIFS server name	i. Disconnect all CIFS access to the 7-Mode system. ii. Use the <code>cifs terminate</code> command to terminate the CIFS service. iii. Use the <code>cifs setup</code> command to create the CIFS server with a different name. The new CIFS server must be created on the 7-Mode system for transitioning CIFS configurations; otherwise, transition of CIFS share-level ACLs fails. The CIFS server name must be unique in the domain. iv. Ensure that the CIFS server is up and running.
Use a new CIFS server name	Ensure that the CIFS server is up and running.

Attention: If you want to preserve the CIFS server name of the 7-Mode system on the Vserver, the number of volumes in the 7-Mode system or vFile unit to be transitioned must be less than or equal to 20 so that all the 7-Mode volumes can be transitioned in a single subproject or session.

- c. Log in as administrator to the cluster.
- d. Use the `system license add` command to add the CIFS license.

Example

```
cluster1::> system license add -license-code AAAAAAAAAAAAAA
```

- e. Verify that at least one data LIF exists on the Vserver.

If a data LIF does not exist, use the `network interface create` command to create the data LIF. For more information about creating data LIFs, see the *Clustered Data ONTAP Network Management Guide*.

- f. Use the `vserver services dns create` command to configure DNS on the Vserver.

Note: If you want to set up the CIFS server on the Vserver by preserving the identity, you must add the same DNS server as that of the 7-Mode system.

Example

```
cluster1::> vserver services dns create -vserver vs_finance -domains  
cifs.lab.example.com -name-servers 192.0.2.145 -state enabled
```

For more information about setting up a DNS domain, see the *Clustered Data ONTAP Network Management Guide*.

- g. Use the `vserver cifs create` command to create a CIFS server on the Vserver.

If you want to preserve the 7-Mode CIFS server name on the Vserver, you must use the original CIFS server name that is used on the 7-Mode system.

For more information about configuring CIFS server, see the *Clustered Data ONTAP File Access Management Guide for CIFS*.

Example

```
cluster1:> vserver cifs create -vserver vs_finance -name CIFSSEVER1 -domain
cifs.lab.example.com
```

8. From the 7-Mode Transition Tool, confirm that the data copy is completed and start the cutover operation.

Note: You must ensure that the secondary session has been completed.

- a. Enter **yes** to start the cutover operation.

Example

```
Step: Start the cutover operation
=====
Do you want to run the complete (cut-over) operation on this session? {yes, no} [yes]:
```

- b. Optional: If you have NFS Kerberos configuration on the 7-Mode system that you want to transition, enter the details of the Kerberos server.

Example

The following example shows a 7-Mode system that has a UNIX-based Kerberos server:

```
NFS Unix Kerberos configured on 7-Mode filer, do you want to transition {yes, no}
[yes]: yes
Enter the FQDN or IP Address of the Key Distribution Center(KDC) Server: 192.0.2.170
Enter the IP address of the Key Distribution Center(KDC) Server
'192.0.2.170' [root]:
Enter the password for 'root@192.0.2.170':
```

The following example shows a 7-Mode system that has an Active Directory based Kerberos server:

```
NFS Microsoft Kerberos configured on 7-Mode filer, do you want to transition {yes, no}
[yes]: yes
Enter the FQDN of the Active Directory Server: AUTH.SEC.EXAMPLE.COM
Enter the IP address of the Active Directory Server 'AUTH.SEC.EXAMPLE.COM'
[192.0.2.171]:
Enter the administrative user-name for the Active Directory Server
'AUTH.SEC.EXAMPLE.COM' [administrator]:
Enter the password for 'administrator@AUTH.SEC.EXAMPLE.COM':
```

Note: The transition of Organizational Unit, in which the machine account of Microsoft Active Directory Kerberos needs to be created, to clustered Data ONTAP 8.2.1 is not supported. As part of transition, the AD Kerberos machine account is created as part of the default Organizational Unit in the AD server.. You should manually modify the Kerberos

configuration for each of the data LIFs transitioned as part of the session to specify the Organizational Unit in which the Microsoft Active Directory Kerberos machine account needs to be created. For more information on how to manage the Kerberos configuration for LIFs, see the *Clustered Data ONTAP File Access Management Guide for NFS*.

- c. Select the mode in which to run the complete operation.

The default mode for completing the transition is `normal`.

Example

This complete operation can be run in either of the following two modes.

Test mode : If you select test mode, all the 7-Mode IP addresses and volumes added to the session are not modified in any way. All the client applications can continue to access data from the 7-Mode storage system after the complete operation. All the LIFs added to the session will be created on the Vserver in the 'down' state.

Normal mode : If you select normal mode, all the 7-Mode IP addresses added to the session will be unconfigured on the 7-Mode storage system and created on the Vserver in the 'up' state. All the 7-Mode volumes added to the session are made offline.

Select the mode in which you want to run this complete operation {normal, test}
[normal]:

- d. Verify the results of the complete operation and rectify any errors or warnings.

9. Verify that the session state is complete by entering **4**.

Example

```
Step: View session status
=====
                Session name : session_purchase2
                Session status : Complete
                Session type  : primary
                Apply configuration : Yes
                7-Mode system FQDN : host1.example.com
                7-Mode version  : NetApp Release 8.1 7-Mode: Fri Nov 16 11:17:42 IST
2012
                vFiler unit   : vfiler0
                Multipath IP address : -
                Cluster-Mode system FQDN : 192.0.2.126
                Cluster-Mode system Version : NetApp Release 8.2.0 Cluster-Mode: Wed Dec 12
01:47:08 IST 2012
                Vserver information : vs39:Pre-existing
                Volume-Pairs       : Purchase:Purchase
                Vserver data LIFs  : -
                SnapMirror schedule : -
                Secondary cluster details : -

7-Mode volume      Cluster-Mode volume      Snapmirror state      Comments
-----
Purchase           Purchase                      -/-                  -
```

A SnapMirror relationship is set up between the clustered Data ONTAP primary volumes to the clustered Data ONTAP secondary volumes. After the complete operation, you must wait till the

SnapMirror resynchronization is successful to make the clustered Data ONTAP secondary volume available for read-only access.

After you finish

Perform any post-transition tasks before enabling client access to the Vserver.

Related concepts

[Preparing for transition](#) on page 12

[Severity levels for precheck messages](#) on page 49

[Guidelines for deciding when to perform cutover](#) on page 47

[Mode types you can use for completing the transition](#) on page 41

[Prerequisites for transitioning NFS and CIFS configurations](#) on page 20

[Error codes and corrective tasks for transition](#) on page 167

Related information

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

Managing transition sessions by using the command-line interface

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

You can manage transition sessions in the following ways:

- Creating a session
- Modifying a session
- Adding information about a cluster, primary 7-Mode system, or secondary cluster to a session
- Managing the credentials of the 7-Mode system and cluster
- Displaying the list of sessions or information about a session
- Verifying the feasibility of transition using the entities in the session
- Starting the data transfer from the 7-Mode volumes to the Clustered Data ONTAP volumes
- Pausing or resuming the data transfer
- Removing information about a primary 7-Mode system or a secondary cluster from a session
- Managing LIFs in transition
- Managing the volumes in transition
- Managing transition jobs
- Completing the transition
- Aborting a transition session
- Deleting a transition session

If any transition operation (such as start, pause, resume, or complete) fails, you must run the transition operation again.

Related concepts

[States of a transition subproject or session](#) on page 42

Related references

[Commands for managing transition](#) on page 143

Creating a transition session by using the command-line interface

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

Before you begin

You must have gathered the following information:

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

Steps

1. From the **Start** menu, click **All Programs > NetApp 7-Mode Transition Tool > NetApp 7-Mode Transition Tool (cli) 1.2** to open the 7-Mode Transition Tool.
The transition main menu is displayed.
2. Exit the transition CLI wizard by entering **4**.

Example

```
Step: Main menu
=====
Select one of the following menu items:
1. Show all sessions
2. Create a new session
3. Select existing session
4. Exit
```

```
Enter your selection(1-4): 4
7-Mode Transition Tool>
```

- From the 7-Mode Transition Tool command prompt, add the user name and password of the systems that are a part of the transition session to the 7-Mode Transition Tool cache by entering the following command:

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

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

user_name is the administrative user name for the system.

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

You are prompted to enter the password of the system.

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

- Create a transition session by entering the following command:

transition create -s *session_name* -t *session_type* -n *7-Mode_system* -c *data-copy-ipaddress* [-f *vfiler_name*] [-h *cluster* -v *vserver_name* -d *schedule* -a {true|false}]

session_name is the name of the transition session.

session_type is the session type. Valid values are standalone, secondary, primary.

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

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

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

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

vserver_name is the name of the Vserver.

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

schedule is the data copy schedule for SnapMirror updates.

When the -a option is set to `true`, all configurations related to the volume, protocol, and services are applied. When set to `false`, only volume configurations are applied and protocol and services configurations are skipped.

Example

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

The transition session moves to the preparation state.

5. Add the 7-Mode volume and clustered Data ONTAP volume to the session by entering the following command:

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

7-Mode_volume is the 7-Mode volume name.

vserver_volume is the clustered Data ONTAP volume on the Vserver.

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

Example

```
7-Mode Transition Tool>transition volumepair add -s session_finance -
v finance -c finance -g aggr1
```

6. Add the IP addresses to be transitioned to the session by entering the following command:

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

Example

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

7. Add the cluster information to the session by entering the following command:

```
transition add-cluster-info -s session_name -h host-name -v vserver-name
[-g aggregate]
```

session_name is the name of the transition session.

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

vserver_name is the name of the Vserver.

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

aggregate is the name of the aggregate on which the root volume of the Vserver is created. The aggregate name is required only if the Vserver has not yet been created.

Example

```
7-Mode Transition Tool>transition add-cluster-info -s session_test -h
10.238.34.241 -v vs_01 -g aggr_01
```

8. For secondary sessions, use the

transition add-primary-seven-mode-system -s *session_name* -h *source-host* -f *management-ipaddress* -d *data-copy-ipaddress* [-m *multipathing-ipaddress*]

command to add information about the 7-Mode primary controller to the transition session. This information is used by the tool to establish cross-mode SnapMirror relationships between the primary 7-Mode controller volumes and the secondary Vserver volumes during the secondary cutover operation. If this information is not provided, the tool will not establish cross-mode SnapMirror relationships as a part of the secondary cutover operation. This command is applicable only to secondary sessions. When transitioning a fan-in relationship, you can add information about more than one 7-Mode primary controller.

session_name is the name of the transition session.

source-host is the IP address or the host name of the primary 7-Mode controllers in the SnapMirror relationships.

Important: The host name must match the value of the *primary-host-or-ip-address* field in the output of the `snapmirror status` command for any of the secondary volumes included in the session.

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

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

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

Example

```
7-Mode Transition Tool>transition add-primary-seven-mode-system -s
sec_session -h 10.238.30.105 -f 10.238.30.107 -d 10.238.30.107
```

9. For primary sessions, use the

transition add-secondary-cluster-mode-system -s *session_name* -h *c-mode-host-name* -v *vserver-name*

command to add information about the cluster to which the secondary volumes in the session have been transitioned. This information is used by the tool to establish cluster-to-cluster SnapMirror relationships between the primary Vserver volumes and the secondary Vserver volumes as a part of the primary cutover. If this information is not provided, the tool will not establish the cluster-to-cluster SnapMirror relationships as a part of the primary cutover operation.

session_name is the name of the transition session.

c-mode-host-name is the FQDN or IP address of the cluster to which the secondary volumes in the session have been transitioned.

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

Example

```
7-Mode Transition Tool>transition add-secondary-cluster-mode-system -
s pri_session -h 10.238.34.241 -v vs_02
```

Running precheck on a session

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

Before you begin

The transition session must be in the preparation state.

About this task

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

Step

1. Verify the compatibility for transition by entering the following command:

`transition precheck -s session_name`

Example

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

[1/14 ] Session status checks
[   Errors   ]

Operation summary:
-----
The 'precheck' operation on the session 'session_finance' has failed.

  1  Errors - Failed!

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

1 Errors:
-----
90202: Checking whether any of the 7-Mode volumes are in 'offline'
```

```

state.                                [      Error      ]
> Following volumes are in 'offline'
state.
> CORRECTIVE-ACTION: Bring offline volumes online by using the following command: 'vol
online
>
<volume_name>'.

> AFFECTED-OBJECTS: { vol2 }

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

Ran precheck for session 'session_finance'.

```

After you finish

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

Related concepts

[Severity levels for precheck messages](#) on page 49

Starting data copy for a transition session

You can initiate a data copy from 7-Mode volumes to clustered Data ONTAP volumes for a baseline transfer. After the baseline transfer is complete, the clustered Data ONTAP volumes are updated periodically based on the SnapMirror schedule selected on the cluster.

Before you begin

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

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

- A SnapMirror schedule for incremental transfers must exist on the cluster.

About this task

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

Step

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

transition start -s session_name

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

Example

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

[1/17 ] Session status checks                      [      Ok      ]
[2/17 ] Validating 7-Mode system information        [      Ok      ]
[3/17 ] Validating 7-Mode volumes information       [      Ok      ]
[4/17 ] Validating Cluster-Mode system information  [      Ok      ]
[5/17 ] Validating Cluster-Mode volume information [      Ok      ]
[6/17 ] Validating LIF information                  [ Notifications ]
[7/17 ] WAFL prechecks                             [      Ok      ]
[8/17 ] UNIX users and groups prechecks            [   Warnings   ]
[9/17 ] SnapMirror prechecks                       [ Notifications ]
[10/17] NFS prechecks                              [   Warnings   ]
[11/17] Networking prechecks                      [ Notifications ]
[12/17] CIFS prechecks                            [   Warnings   ]
[13/17] Name services prechecks                   [   Warnings   ]
[14/17] SAN prechecks                             [      Ok      ]
[15/17] Creating Cluster-Mode volumes              [      Ok      ]
[16/17] Establishing SnapMirror relationships between the 7-Mode and Cluster-Mode
volumes                                           [      Ok      ]
[17/17] Initializing SnapMirror relationships between the 7-Mode and Cluster-Mode
volumes                                           [      Ok      ]

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

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

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

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

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

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

Setting the audit log path of the Vserver

To enable the tool for transitioning the audit configuration, you must provide the path for saving the audit logs on the Vserver by using the 7-Mode Transition Tool CLI.

About this task

If auditing is enabled on the 7-Mode storage system and you have not provided the path for saving the audit logs on the Vserver, then the precheck operation displays a warning message asking you to provide the details.

Step

1. To set the path for saving the audit logs on the Vserver, provide the audit path to the session by entering the following command:

```
transition property-set -s <session_name> -p cluster-audit-saveas-path -v <existing_path_on_vserver> session_name
```

The transition complete operation uses the specified path for saving the audit log files while transitioning the audit configuration from the 7-Mode storage system to the destination Vserver.

Note: For a new Vserver created during the transition start operation, you can set the property only after the transition start operation finishes (that is, after the Vserver is created). You can set and modify this property even after the preparation phase.

Pausing and resuming a copy operation

You can pause or stop all ongoing and scheduled SnapMirror transfers such as those occurring during a controller takeover or employed for troubleshooting. You can then resume the data copy of a paused session.

Before you begin

- To pause a session that is in the copy-baseline state, all the volumes in the session must have a restart checkpoint available.

About this task

When you pause a primary session, only the SnapMirror transfers between the 7-Mode and clustered Data ONTAP primary volumes are paused. The SnapMirror transfers between the 7-Mode primary and clustered Data ONTAP secondary volumes are not paused.

Steps

1. To pause a session, enter the following command:

```
transition pause -s session_name
```

The session moves to the copy-pause state.

2. To resume the data copy of a paused session, enter the following command:

```
transition resume -s session_name
```

If you paused the session when transition was in the copy-baseline state, the resumed session returns to the copy-baseline state. Similarly, if the session was paused at copy-update, the resumed session returns to the copy-update state.

Completing the transition

After you have verified and decided to cutover, you can complete the transition. You can complete the transition in either normal mode or test mode.

Before you begin

If you are running transition complete in normal mode, you must have disconnected the clients that are accessing the 7-Mode volumes.

Attention: Because transition cutover is disruptive to the clients, you must plan the activity to minimize the downtime.

About this task

The transition complete operation might take several minutes to complete, depending on the number of volumes in the session, the time taken for the final data transfer, and the configurations to be applied on the cluster.

Step

1. Complete the transition process by entering the following command:

```
transition complete -s session_name -m
```

session_name is the name of the session.

-m: Completes the transition in test mode. If the complete operation is run in test mode, the 7-Mode volumes are not taken offline and IP addresses are not removed. Clients can continue to access the 7-Mode volumes. All configurations are applied to the clustered Data ONTAP volumes, and the data LIFs are configured on the Vserver in the down state.

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

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

After you finish

Once the session state changes to completed, you cannot re-run the `transition complete` command. Therefore, if you complete the transition in test mode and want to enable client access to

the clustered Data ONTAP volumes, you must manually take the 7-Mode volumes offline, remove the 7-Mode IP addresses, and configure the data LIFs on the Vserver to the up state.

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

Related concepts

[Mode types you can use for completing the transition](#) on page 41

Displaying transition session details

You can display transition session information, such as the session name, session state, 7-Mode system FQDN, Data ONTAP version, Vserver information, and SnapMirror schedule. After baseline copy is completed and a session moves to the copy-update state, you must monitor the state of the session by viewing the session status and prepare for cutover.

Step

- 1. View the session status by entering the following command:

```
transition show -s session_name -c
```

Example

The following example shows the details of the session session_finance:

```
7-Mode Transition Tool>transition show -s session_finance -c
      Session name : session_finance
      Session status : Complete
      Session type : standalone
      Apply configuration : Yes
      7-Mode system FQDN : myhost.example.com
      7-Mode version : NetApp Release 8.1 7-Mode: Fri Nov 16
11:17:42 IST 2012
      vFiler unit : vfiler0
      Multipath IP address : -
      Cluster-Mode system FQDN : 192.0.2.149
      Cluster-Mode system Version : NetApp Release 8.2.0 Cluster-Mode: Wed
Dec 12 01:47:08 IST 2012
      Vserver information : vs49:Pre-existing
      Volume-Pairs : testVol1:testVol1
      Vserver data LIFs : -
      SnapMirror schedule : daily

7-Mode volume      Cluster-Mode volume      Snapmirror state      Comments
-----
testVol1           testVol1                 -/-                   -
```

Aborting a transition session

You can abort a transition session if there are issues during the data copy phase. When a session is aborted, all data transfers between the 7- Mode volumes and clustered Data ONTAP volumes are stopped and all SnapMirror relationships between the 7-Mode system and the Vserver are deleted.

Before you begin

The session must be in copy-baseline, copy-update, copy-pause, or cutover state.

Step

1. To abort a session, enter the following command:

```
transition abort -s session_name
```

The only operation that you can perform on an Aborted session is to delete it.

Deleting a transition session

You can delete a specified transition session. The session can be deleted only if it is in the preparation, completed, or aborted state.

Before you begin

A session that is in the copy-baseline, copy-update, or copy-pause state must be aborted before it can be deleted.

Step

1. To delete a session, enter the following command:

```
transition delete -s session_name
```

Example

```
7-Mode Transition Tool>transition delete -s session_vol0
```

Viewing transition jobs

Every time a transition operation is executed, a transition job is created by the 7-Mode Transition Tool. You can view all the jobs of a session. You can also view the status and result of each job to view the errors and warnings, if any, for each operation. You can also redirect the output of a job result to a text file to view the results offline.

About this task

Jobs are created for the following transition operations:

- transition precheck
- transition start
- transition pause
- transition resume
- transition abort
- transition complete

Step

1. Depending on the job details that you want to view, perform the appropriate action:

If you want to...	Enter this command...
-------------------	-----------------------

View the jobs in a session

transition jobs -s *session_name*

The following example shows the jobs of the session named session_finance:

```
7-Mode Transition Tool>transition jobs -s session_finance
time      Unique job identifier      Operation      Status      Start
-----
-----
caa97f28-5cfd-4882-b18b-b80bf36307b7  complete  completed  12-12-2012 07:01:01
EST 12-12-2012 07:01:22 EST
2bd46417-0e0f-4330-b296-4e19e4b610d7  start     completed  12-12-2012 06:53:43
EST 12-12-2012 06:53:54 EST
03bcd690-2e85-41e6-a46b-4cfde54a6202  start     completed  12-12-2012 06:53:11
EST 12-12-2012 06:53:21 EST
9b9e6e73-aa49-4eee-bbf3-321d820d0730  start     completed  12-12-2012 06:51:20
EST 12-12-2012 06:51:30 EST
44740214-1cd0-4d55-aa99-5a78603e9c22  precheck  completed  12-12-2012 06:50:04
EST 12-12-2012 06:50:13 EST
8dcba885-f1bd-496e-ab59-0c10e46af60d  precheck  completed  12-12-2012 06:18:50
EST 12-12-2012 06:18:57 EST
```

View the job status

transition job-status -j *job_id*

The job status can be one of the possible values: running, completed, rejected, or aborted.

The following example shows the status of the start operation for a particular session:

```
7-Mode Transition Tool>transition job-status -j 02a86aba-34fe-418a-881a-51097b0c0cf2
Job id : 02a86aba-34fe-418a-881a-51097b0c0cf2
Job type : start
Job status : completed
```

If you want to...	Enter this command...
--------------------------	------------------------------

View the job results	transition job-results -j <i>job_id</i>
----------------------	--

The following example shows the results of the start operation for a particular session:

```
7-Mode Transition Tool>transition job-results -j 02a86aba-34fe-418a-881a-51097b0c0cf2
[1/17 ] Session status checks [ Ok ]
[2/17 ] Validating 7-Mode system information [ Ok ]
[3/17 ] Validating 7-Mode volumes information [ Ok ]
[4/17 ] Validating Cluster-Mode system information [ Ok ]
[5/17 ] Validating Cluster-Mode volume information [ Ok ]
[6/17 ] Validating LIF information [ Ok ]
[7/17 ] WAFL prechecks [ Ok ]
[8/17 ] UNIX users and groups prechecks [ Ok ]
[9/17 ] SnapMirror prechecks [ Warnings ]
[10/17 ] NFS prechecks [ Warnings ]
[11/17 ] Networking prechecks [ Notifications ]
[12/17 ] CIFS prechecks [ Ok ]
[13/17 ] Name services prechecks [ Warnings ]
[14/17 ] SAN prechecks [ Ok ]
[15/17 ] Creating Cluster-Mode volumes [ Ok ]
[16/17 ] Establishing SnapMirror relationships between the 7-Mode and Cluster-Mode
volumes [ Ok ]
[17/17 ] Initializing SnapMirror relationships between the 7-Mode and Cluster-Mode
volumes [ Ok ]

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

 0 Errors - No action required
11 Warnings - Need your attention!
13 Notifications - Please review
81 Informational messages

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

Export the job results to a text file	transition job-results -j <i>job_id</i> -r no>filepath
---------------------------------------	--

The following example shows how to redirect the job results to a file named `transition_job_results`:

```
transition job-results -j
02a86aba-34fe-418a-881a-51097b0c0cf2 -r no>C:
\transition_job_result.rtf
```

Commands for managing transition

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

- [Commands to manage transition sessions](#) on page 144
- [Commands to manage transition operations](#) on page 144
- [Commands to manage volume SnapMirror relationships](#) on page 144
- [Commands to manage credentials of Data ONTAP systems](#) on page 145
- [Commands to manage the volume pairs in a session](#) on page 145
- [Commands to manage LIF details](#) on page 145

- [Commands to manage transition jobs](#) on page 146

Commands to manage transition sessions

If you want to...	Use this command...
Run the end-to-end transition wizard	<code>transition menu</code>
Modify a transition session	<code>transition modify</code>

Commands to manage transition operations

If you want to...	Use this command...
Create a transition session	<code>transition create</code>
Start data copy for the volumes in a session	<code>transition start</code>
Pause data copy for the volumes in a session	<code>transition pause</code>
Resume data copy of a paused session	<code>transition resume</code>
Abort a session	<code>transition abort</code>
Delete a session	<code>transition delete</code>
Display the details of a transition session	<code>transition show</code>

Commands to manage volume SnapMirror relationships

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

If you want to...	Use this command...
View information about the clustered Data ONTAP secondary systems added to a primary session	<code>transition show-secondary-cluster-mode-systems</code>

Commands to manage credentials of Data ONTAP systems

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

Commands to manage the volume pairs in a session

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

Commands to manage LIF details

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

Commands to manage session properties

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

Commands to manage transition jobs

If you want to...	Use this command...
Display the list of jobs in a session	<code>transition jobs</code>
View the status of a job	<code>transition job-status</code>
View the results of a job	<code>transition job-results</code>

The following option can be used with all commands:

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

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

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

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.

Log files for the 7-Mode Transition Tool

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

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

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

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

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

Log file path	Contains information about...
<code>transition-cli.log</code>	Debug messages that are part of the CLI wizard
<code>transition-gui.log</code>	Entries of all the actions performed by using the web interface
<code>zapi-outbound-cli.log</code>	All the Data ONTAP APIs and their outputs that are executed by the 7-Mode Transition Tool when the CLI wizard is run
<code>default/audit.log</code>	- All the parameters—such as HTTP or HTTPS port and log directory path—that are used by the tool every time the 7-Mode Transition Tool is run - All the transition commands that are executed with the outputs
<code>default/transition.log</code>	Debug messages that are not specific to any session
<code>default/zapi-outbound.log</code>	Output of all the Data ONTAP APIs that are executed by the 7-Mode Transition Tool and which do not belong to any session
<code>server-console.log</code>	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.

Error messages in 7-Mode Transition Tool

You might encounter issues when using the 7-Mode Transition Tool. The tool displays the error messages and possible solutions. You can also use the log files of the tool and the transition session files to troubleshoot the issues.

Failed to create the specified directory

Message Creating a directory '/vol' on the root volume of the Vserver 'NSTFAS32704VS4'. This is required to create junction paths for the Vserver volumes. Failed to create the specified directory.

Reason: Read-only file system.

CORRECTIVE-ACTION: Create junction paths for the Vserver volumes manually after the 'transition complete' operation using the following command: 'volume mount -vserver NSTFAS32704VS4 -volume <volume_name> -junction-path <junction_path>'

For more information, refer to Data and Configuration Transition guide.

- Cause** The root volume has LS mirrors. All the LS mirror volumes are read-only by default. This issue occurs when the tool tries to create a /vol directory on the root volume, and the request lands on a node that has root volume LS mirror.
- Corrective action** Delete the LS mirrors of the root volume, and then try the operation again.

Resolving BUILTIN CIFS local group name conflicts

- Message** The following BUILTIN CIFS local group entries have name conflicts and cannot be transitioned to the Vserver 'vs_standalone_15'.
- CORRECTIVE-ACTION: Rename the BUILTIN CIFS local groups on the Vserver to remove the conflicts, and then retry the operation.
- For more information, refer to Data and Configuration Transition guide.
- AFFECTED-OBJECTS: { 7-Mode : BUILTIN\GroupU -> Cluster : BUILTIN\GroupA; 7-Mode :BUILTIN\GroupA -> Cluster : BUILTIN\GroupU }
- Cause** The BUILTIN CIFS local group entries have name conflicts.
- Corrective action** Manually resolve the conflicts with CIFS BUILTIN local group names by renaming one of the groups that has the group name conflict to a temporary name that is not in use on the cluster, and then retry the operation. For example, rename either of the groups BUILTIN\GroupA, BUILTIN\GroupU to a username that is not being used by any of the local users on the Vserver.
- Another example would be: cifs users-and-groups local-user rename -user-name BUILTIN\GroupA -new-user-name BUILTIN\GroupX -vserver vs_standalone_15

Failed to communicate with 7-Mode Transition Tool server

- Message** Failed to communicate with 7-Mode Transition Tool server.
- Cause** The NetApp Transition Tool service is not running.

- Corrective action**
1. Open Windows Task Manager and click the **Services** tab.
 2. Click **Services**.
 3. Right-click the **NetApp Transition Tool** and select **Start**.

Host not found in credential bank

Message Host *host_name* not found in credential bank.

- Cause**
- The NetApp Transition Tool service was restarted and the credentials were lost.
 - Credentials were removed by using the CLI.

- Corrective action**
1. From the 7-Mode Transition Tool CLI, use the `transition credentials get` command to verify that the credentials are not available on the transition server.

```
7-Mode Transition Tool>transition credentials get -h 192.238.55.34
Error(21500) : Host '192.238.55.34' not found in credential bank.
```

2. Use the `transition credentials add` command to add the credentials of the cluster.

```
7-Mode Transition Tool>transition credentials add -h 192.238.55.34 -u
admin
Enter password for 'admin@192.238.55.34':
```

Error messages when opening 7-Mode Transition Tool (CLI)

Details Immediately after launching the 7-Mode Transition Tool CLI wizard, you might see a series of error messages.

An output similar to the following is displayed:

```
*****
*                               Welcome to NetApp 7-Mode Transition Tool.                               *
*                               *                                       *                               *
* - Use 'transition menu' to perform all the steps of transition process.                               *
* - Type 'help' or '?' to see the list of commands provided by the tool.                               *
* - To get help on any command use 'transition <command name> help'.                               *
* Example: 'transition menu help'                                       *
* - Recommended terminal width is 120 or more characters for best experience.                       *
* Default terminal width is 125 characters.                             *
*****
16:10:43,054 |-INFO in ch.qos.logback.classic.LoggerContext[default] - Could NOT find
resource [logback.groovy]
16:10:43,055 |-INFO in ch.qos.logback.classic.LoggerContext[default] - Could NOT find
resource [logback-test.xml]
16:10:43,056 |-INFO in ch.qos.logback.classic.LoggerContext[default] - Found resource
[logback.xml] at [file:/C:/Program%20Fi
les/NetApp/7-Mode%20Transition%20Tool/etc/logback.xml]
16:10:43,249 |-INFO in ch.qos.logback.classic.joran.action.ConfigurationAction - debug
attribute not set
16:10:43,309 |-INFO in ch.qos.logback.core.joran.action.AppenderAction - About to
instantiate appender of type [ch.qos.logbac
k.core.rolling.RollingFileAppender]
```

```
16:10:43,317 |-INFO in ch.qos.logback.core.joran.action.AppenderAction - Naming appender as [audit]
16:10:43,387 |-INFO in ch.qos.logback.core.joran.action.NestedComplexPropertyIA - Assuming default type [ch.qos.logback.classic.encoder.PatternLayoutEncoder] for [encoder] property
16:10:43,596 |-INFO in ch.qos.logback.core.rolling.FixedWindowRollingPolicy@214a55f2 - Will use zip compression
16:10:43,623 |-INFO in ch.qos.logback.core.rolling.RollingFileAppender[audit] - Active log file name: logs/audit.log
16:10:43,623 |-INFO in ch.qos.logback.core.rolling.RollingFileAppender[audit] - File property is set to [logs/audit.log]
16:10:43,628 |-ERROR in ch.qos.logback.core.rolling.RollingFileAppender[audit] - openFile(logs/audit.log,true) call failed. j
ava.io.FileNotFoundException: logs\audit.log (Access is denied)
        at java.io.FileNotFoundException: logs\audit.log (Access is denied)
        at java.io.FileOutputStream.openAppend(Native Method)
        at java.io.FileOutputStream.<init>(Unknown Source)
```

- Cause** You have logged in to Windows system as a nonadministrator user.
- Corrective action** You must have administrator privileges on the Windows system to run the 7-Mode Transition Tool.

Failed to fetch volume information from the 7-Mode system

- Message** Failed to fetch volume information from the 7-Mode system
- Cause** Some operations might be in progress on the 7-Mode volumes when the tool is collecting 7-Mode volume information. These operations can be an aggregate or a volume going offline, a volume being deleted, and so on.
- Corrective action** Retry the operation by selecting volume configuration from the Setup menu.

Error occurred with Cluster-Mode storage system

- Message** ERROR: Error occurred with Cluster-mode storage system *cluster*. Error Message :RPC: Port mapper failure - RPC: Timed out.
- Cause** One of the nodes in the cluster is in takeover mode and is waiting for a giveback.
- Corrective action** Wait till the giveback operation is complete and all the nodes of the cluster are in healthy state.

Related concepts

[Impact of takeover and giveback on transition](#) on page 176

Error onlining volume

- Error message** Error onlining volume "vserver_name:volume". Unable to set volume attributes for volume "Vserver_volume" on Vserver "Vserver". Reason: Volume was left in an inconsistent state by an aborted copy or an aborted snapmirror initial (level 0)

transfer. In order to bring it online, you must either destroy and recreate the volume, or complete an initial snapmirror transfer or copy.

Cause The clustered Data ONTAP volume was left in an inconsistent state because the SnapMirror initialization was aborted.

Corrective action

1. Log in to the cluster.
2. Use the `snapmirror show` command to verify the SnapMirror status of the volume. The following example shows how to identify an aborted SnapMirror initialization:

```
cluster1::> snapmirror show -destination-path vs1:vol_dst
Source Path: vs2:vol_src
Destination Path: vs1:vol_dst
Relationship Type: TDP
SnapMirror Schedule: -
Tries Limit: -
Throttle (KB/sec): unlimited
Mirror State: Uninitialized
Relationship Status: Idle
Transfer Snapshot: -
Snapshot Progress: -
Total Progress: -
Snapshot Checkpoint: -
Newest Snapshot: -
Newest Snapshot Timestamp: -
Exported Snapshot: -
Exported Snapshot Timestamp: -
Healthy: false
Unhealthy Reason: Transfer failed.
Constituent Relationship: false
Destination Volume Node: cluster1-01
Relationship ID: 59fcdaf2-4b26-11e2-a038-123478563412
Transfer Type: -
Transfer Error: -
Current Throttle: -
Current Transfer Priority: -
Last Transfer Type: initialize
Last Transfer Error:
Transfer aborted.
Last Transfer Size: -
Last Transfer Duration: -
Last Transfer From: vs2:vol_src
Last Transfer End Timestamp: 12/21 04:25:55
Progress Last Updated: -
Relationship Capability: 8.2 and above
Lag Time: -
SnapMirror Policy: DPDefault
```

The Snapshot Checkpoint field is empty, the Last Transfer Type is initialize, and the Last Transfer Error is Transfer aborted, indicating that the volume is in an inconsistent state and cannot be recovered by running the transition start operation.

3. Use the `snapmirror break` command to break the SnapMirror relationship of this volume.

```
cluster1::> snapmirror break -destination-path vs1:vol_dst
```

4. Use the `snapmirror delete` command to delete the SnapMirror relationship of this volume.

```
cluster1::> snapmirror delete -destination-path
vs1:vol_dst
```

5. Use the `volume offline` command to take the volume offline.

```
cluster1::> volume offline -vserver vs1 -volume vol_dst
```

6. Use the `volume create` command to create a new volume of type DP with the same name and properties as the deleted volume.

```
cluster1::> volume create -vserver vs1 -volume vol_dst -
aggregate aggr1 -size 50m -state online -type DP
```

7. From the 7-Mode Transition Tool, re-run the transition start operation.

Network connectivity issues with the 7-Mode system

Details

The following example outputs show errors related to connectivity problems:

```
7-Mode Transition Tool>transition credentials add -h 192.238.55.34 -u root
Enter password for 'root@192.238.55.34':
Error(31010) : Host '192.238.55.34' is not reachable.
```

```
01032: Communication problem occurred.
> Timeout occurred while communicating with 7-Mode storage system
'192.238.55.34'.
> CORRECTIVE-ACTION: Check your network configuration, correct the problem
and re-run the operation
```

```
Step: Select the 7-Mode storage system
=====
You can type 'back', 'exit', 'help' or '?' at any time.

Enter the FQDN or IP address of the 7-Mode system [192.238.55.34]:
Enter the administrative user-name for the 7-Mode system '192.238.55.34'
[root]:
Enter the password for 'root@192.238.55.34':
ERROR: Host '192.238.55.34' timed out.
```

Cause 1

The 7-Mode system is down or unreachable.

Corrective action 1

1. From the command prompt of the Windows machine, ping the 7-Mode system by using the following command:

ping *IP_address*

2. If the 7-Mode system cannot be pinged, resolve any network issues and make the 7-Mode system reachable.

Cause 2 HTTP is disabled on the 7-Mode system, but the 7-Mode Transition Tool uses HTTP to communicate with the 7-Mode system.

Corrective action 2

1. Open the `\etc\conf\transition-tool.conf` file to verify that the tool uses HTTP to communicate with the 7-Mode system.

The following example shows the contents of such a file:

```
#####
# Data ONTAP API HTTP/HTTPS configurations
#####

# Default value is false
ontap.http.enabled = true

# Default value is true
ontap.https.enabled = false
```

2. Log in to the 7-Mode system.
3. Enable HTTP on the 7-Mode system by entering the following command:
options httpd.admin.enable on

Cause 3 SSLv3 is disabled on the 7-Mode system, but the tool is using HTTPS (through SSLv3) to communicate with the 7-Mode system

Corrective action 3

1. Open the `\etc\conf\transition-tool.conf` file to verify that the tool uses HTTPS to communicate with the 7-Mode system.

The following example shows the contents of such a file:

```
#####
# Data ONTAP API HTTP/HTTPS configurations
#####

# Default value is false
ontap.http.enabled = false

# Default value is true
ontap.https.enabled = true
```

2. Log in to the 7-Mode system.
3. Enable SSLv3 on the 7-Mode system by entering the following command:
options ssl.v3.enable on
4. Enable HTTPS on the 7-Mode system by entering the following command:

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

Related tasks

Preparing the 7-Mode system for transition on page 18

Network connectivity issues with the cluster

Details The following example outputs show errors related to connectivity problems:

```
7-Mode Transition Tool>transition credentials add -h 192.238.55.35 -u admin
Enter password for 'admin@192.238.55.35':
Error(31009) : Host '192.238.55.35' timed out.
```

```
01032: Communication problem
occurred.
[      Error      ]
> Timeout occurred while communicating with Cluster-Mode storage system
'192.238.55.35'.
> CORRECTIVE-ACTION: Check your network configuration, correct the problem and
re-run the
> operation.
```

```
Step: Select target cluster
=====
Enter the FQDN or management IP address of the cluster: 192.238.55.35
Enter the administrative user-name for the cluster '192.238.55.35' [admin]:
Enter the password for 'admin@10.238.55.32':
ERROR: Host '192.238.55.35' timed out.
```

Cause 1 The cluster is down or unreachable.

Corrective action 1 1. From the command prompt of the Windows machine, ping the cluster by using the following command:

```
ping cluster-mgmt_IP_address
```

2. If there is no response to the ping request from the cluster, bring the IP address to the up state, and try again.

Cause 2 HTTP is disabled on the cluster, but the 7-Mode Transition Tool uses HTTP to communicate with the cluster.

Corrective action 2 1. Open the `\etc\conf\transition-tool.conf` file to verify that the tool uses HTTP to communicate with the cluster.

The following example shows the contents of such a file:

```
#####
# Data ONTAP API HTTP/HTTPS configurations
#####
```

```
# Default value is false
ontap.http.enabled = true
```

```
# Default value is true
ontap.https.enabled = false
```

2. Log in to the cluster.
3. Use the `network interface show` command to view the firewall policy of the LIF that is used to communicate with the 7-Mode Transition Tool.

```
cluster1::> network interface show -vserver vs1 -lif cluster_mgmt
      Vserver Name: vs1
Logical Interface Name: cluster_mgmt
      Role: cluster-mgmt
      Data Protocol: none
      Home Node: node1
      Home Port: e0c
      Current Node: node1
      Current Port: e0c
Operational Status: up
Extended Status: -
      Is Home: true
Network Address: 192.0.2.134
      Netmask: 255.255.255.192
      Bits in the Netmask: 18
      IPv4 Link Local: -
      Routing Group Name: c192.0.0.0/18
Administrative Status: up
      Failover Policy: nextavail
      Firewall Policy: policy123
      Auto Revert: false
Fully Qualified DNS Zone Name: none
      DNS Query Listen Enable: false
      Failover Group Name: system-defined
      FCP WPN: -
      Address family: ipv4
      Comment: -
```

4. Use the `system services firewall policy show` command to determine whether the firewall policy associated with the cluster-management LIF is blocking the HTTP service.

```
cluster1::> system services firewall policy show -policy policy123
Policy      Service      Action IP-List
-----
policy123
      dns      allow  0.0.0.0/0, ::/0
      http      deny  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.
```

5. Use the `system services firewall policy create` command to create a firewall policy that allows HTTP service along with other services that are required.

```
cluster1::> system services firewall policy create -policy
policyAllow -service http,service1,service2 -ip-list
0.0.0.0/0 -action allow
```

Specify the other services that must be allowed for the policy as a comma-separated list.

6. Use the `network interface modify` command to apply the firewall policy to the LIF that is used to communicate with the 7-Mode Transition Tool.

```
cluster1::> network interface modify -vserver vs1 -lif
cluster_mgmt -firewall-policy policyAllow
```

Cause 3 SSLv3 is disabled on the cluster, but the tool is using HTTPS (through SSLv3) to communicate with cluster.

Corrective action 3 1. Open the `\etc\conf\transition-tool.conf` file to verify that the tool uses HTTPS to communicate with the cluster.

The following example shows the contents of such a file:

```
#####
# Data ONTAP API HTTP/HTTPS configurations
#####
```

```
# Default value is false
ontap.http.enabled = false
```

```
# Default value is true
ontap.https.enabled = true
```

2. Log in to the cluster.
3. Use the `system services web show` command to view whether SSLv3 is disabled on the cluster.

```
cluster1::> system services web show
External Web Services: true
Status: online
HTTP Protocol Port: 80
HTTPS Protocol Port: 443
TLsv1 Enabled: true
SSLv3 Enabled: false
SSLv2 Enabled: true
```

Cause 4 HTTPS is disallowed on the cluster, but the tool is using HTTPS to communicate with the cluster.

Corrective action 4 1. Open the `\etc\conf\transition-tool.conf` file to verify that the tool uses HTTP to communicate with the cluster.

The following example shows the contents of such a file:

```
#####
# Data ONTAP API HTTP/HTTPS configurations
#####
```

```
# Default value is false
ontap.http.enabled = false
```

```
# Default value is true
ontap.https.enabled = true
```

2. Log in to the cluster.
3. Use the `network interface show` command to view the firewall policy of the LIF that is used to communicate with the 7-Mode Transition Tool.

```
cluster1::> network interface show -vservers vs1 -lif cluster_mgmt
      Vserver Name: vs1
      Logical Interface Name: cluster_mgmt
      Role: cluster-mgmt
      Data Protocol: none
      Home Node: node1
      Home Port: e0c
      Current Node: node1
      Current Port: e0c
      Operational Status: up
      Extended Status: -
      Is Home: true
      Network Address: 192.0.2.134
      Netmask: 255.255.255.192
      Bits in the Netmask: 18
      IPv4 Link Local: -
      Routing Group Name: c192.0.0.0/18
      Administrative Status: up
      Failover Policy: nextavail
      Firewall Policy: policy123
      Auto Revert: false
      Fully Qualified DNS Zone Name: none
      DNS Query Listen Enable: false
      Failover Group Name: system-defined
      FCP WWP: -
      Address family: ipv4
      Comment: -
```

4. Use the `system services firewall policy show` command to determine whether the firewall policy associated with the cluster-management LIF is blocking the HTTPS service.

```
cluster1::> system services firewall policy show -policy policy123
Policy      Service      Action IP-List
```

```
-----
policy123
      dns      allow  0.0.0.0/0, ::/0
      http     allow  0.0.0.0/0, ::/0
      https    deny   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.
```

5. Use the system services firewall policy create command to create a firewall policy that allows HTTPS service along with other services that are required.

```
cluster1::> system services firewall policy create -policy
policyAllow -service https,service1,service2 -ip-list
0.0.0.0/0 -action allow
```

Specify the other services that must be allowed for the policy as a comma-separated list.

6. Use the network interface modify command to apply the firewall policy to the LIF that is used to communicate with the 7-Mode Transition Tool.

```
cluster1::> network interface modify -vserver vs1 -lif
cluster_mgmt -firewall-policy policyAllow
```

Related tasks

[Preparing the cluster for transition](#) on page 15

Windows system out of memory

Details The following example outputs show errors related to low system memory:

```
....
Unexpected error occurred while detecting java. Error code: 1.
Close other programs. And try again later.
Press any key to continue . . .
```

```
...
Unexpected error occurred while detecting java. Error code: 1.
Close other programs. And try again later.
Press any key to continue . . .
```

```
Step: Select the 7-Mode storage system
=====
You can type 'back', 'exit', 'help' or '?' at any time.
Enter the FQDN or IP address of the 7-Mode system: 192.72.219.21
Enter the administrative user-name for the 7-Mode system '192.72.219.21' [root]:
Enter the password for 'root@10.72.219.21':
Error: null
```

```
Exception in thread "main" java.lang.OutOfMemoryError
at java.util.zip.Inflater.inflateBytes(Native Method)
at java.util.zip.Inflater.inflate(Unknown Source)
at java.util.zip.InflaterInputStream.read(Unknown Source)
at sun.misc.Resource.getBytes(Unknown Source)
at java.net.URLClassLoader.defineClass(Unknown Source)
at java.net.URLClassLoader.access$100(Unknown Source)
at java.net.URLClassLoader$1.run(Unknown Source)
at java.net.URLClassLoader$1.run(Unknown Source)
at java.security.AccessController.doPrivileged(Native Method)
at java.net.URLClassLoader.findClass(Unknown Source)
at java.lang.ClassLoader.loadClass(Unknown Source)
at sun.misc.Launcher$AppClassLoader.loadClass(Unknown Source)
at java.lang.ClassLoader.loadClass(Unknown Source)
at com.netapp.transition.cli.logging.CliAuditLogger.logMessage(CliAuditLogger.java:99)
at com.netapp.transition.cli.logging.CliAuditLogger.error(CliAuditLogger.java:74)
at com.netapp.transition.cli.logging.AuditLoggingCliResponseListener.handleResponse
(AuditLoggingCliResponseListener.java:77)
at com.netapp.autocli.CliExecutor.execute(CliExecutor.java:279)
at com.netapp.transition.cli.CliMain.main(CliMain.java:64)
```

Cause High memory usage by other processes running on the Windows workstation.

- Corrective action**
1. Open Windows Task Manager and click the **Performance** tab.
 2. In the Physical Memory (MB) area, verify that the free memory is low (usually less than 50 MB).
 3. Click the **Processes** tab.
 4. Select the unused processes and click **End Process**.
 5. Rerun the 7-Mode Transition Tool.

Related information

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

Session state changes to copy-unknown

Issue Session state changes to copy-unknown.

Cause 1 Credentials of the cluster are lost from the transition server because either the server is restarted or the credentials are removed from the server.

You can view the session status to view the state. The following output shows the error message that is displayed for this issue:

```
Session name : Session_20130220_101147
Session status : Copy-unknown
Session type : standalone
Apply configuration : Yes
7-Mode system FQDN : 192.238.55.28
7-Mode version : xxx
vFiler unit : vfiler0
Multipath IP address : -
Cluster-Mode system FQDN : 192.238.55.34
Cluster-Mode system Version : xxx
Vserver information : vs1:Pre-existing
Volume-Pairs : vol0:vol0
Vserver data LIFs : -
SnapMirror schedule : daily
```

7-Mode volume	Cluster-Mode volume	Snapmirror state	Comments
vol0	vol0	-/-	* Error: Host
'192.238.55.34' not found in credential bank.			

- Corrective action for cause 1**
1. From the 7-Mode Transition Tool CLI, use the `transition credentials get` command to verify that the credentials are not available on the transition server.

```
7-Mode Transition Tool>transition credentials get -h 192.238.55.34
Error(21500) : Host '192.238.55.34' not found in credential bank.
```

2. Use the `transition credentials add` command to add the credentials of the cluster.

```
7-Mode Transition Tool>transition credentials add -h 192.238.55.34 -u admin
Enter password for 'admin@192.238.55.34':
```

- Cause 2** Transition start operation failed on some volumes in the session.

The following example shows the part of the session status output that indicates this issue:

```
.....
7-Mode volume Cluster-Mode volume  Snapmirror state
Comments
-----
vol1          vol1                uninitialized/
transferring
-
vol2          vol2_21              -/-                      * Error:
SnapMirror relationship does not exist.
vol3          vol3                -/-                      * Error:
SnapMirror relationship does not exist.
```

- Corrective action for cause 2**
1. Verify the results of the start operation and resolve any blocking errors.
 2. Use the `transition start` command to rerun the start operation.

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

3. Use the `transition show` command to verify that the start operation was successful on all the volumes in the session.

Related concepts

[States of a transition subproject or session](#) on page 42

Final update failed for one or more volumes during transition complete

Message	Final update failed for the clustered Data ONTAP volume <code>volume_name</code> . Reason: Transfer failed.
Cause 1	The 7-Mode volume contains the maximum limit of 255 Snapshot copies and cannot create a Snapshot copy for the final update transfer.
Corrective action 1	Delete one or more Snapshot copies of the 7-Mode volume on page 161
Cause 2	The latest common Snapshot copy between the 7-Mode system and the Vserver is deleted from the 7-Mode system.
Corrective action 2	Restore the latest common Snapshot copy between the clustered Data ONTAP volume and the 7-Mode volume on the 7-Mode system on page 162
Cause 3	Volume does not have enough space to create a Snapshot copy for the final update transfer.
Corrective action 3	Ensure that the volume has sufficient space to create Snapshot copies. For more information about space allocation in volumes, see the <i>Clustered Data ONTAP Logical Storage Management Guide</i> .

Deleting a Snapshot copy from the 7-Mode volume

If the final update transfer for any volume fails during the transition complete operation because the number of Snapshot copies on the 7-Mode system has exceeded 255 copies, you must delete one or more Snapshot copies.

About this task

The complete operation fails with the following message: Final update failed for the clustered Data ONTAP volume `volume_name`. Reason: Transfer failed.

Steps

1. Log in to the 7-Mode system.
2. View the total number of Snapshot copies for the 7-Mode volume by entering the following command:
snap list -v volume_name
3. If the total number of Snapshot copies is 255, delete one or more Snapshot copies for the volume by entering the following command:
snap delete -v volume_name snapshot_name
4. Log in to the 7-Mode Transition Tool.

5. Run the transition complete operation by entering the following command:

```
transition complete -s session_name [-m]
```

Restoring the latest common Snapshot copy on the 7-Mode system

If the final update transfer for any volume fails during the transition complete operation because the latest common Snapshot copy between the 7-Mode system and the Vserver is deleted from the 7-Mode system, you must restore the latest Snapshot copy on the 7-Mode system.

About this task

The complete operation fails with the following message: Final update failed for the clustered Data ONTAP volume *volume_name*. Reason: Transfer failed.

Steps

1. Verify what is the latest Snapshot copy of the clustered Data ONTAP volume.
 - a. Log in to the cluster.
 - b. Verify what is the latest Snapshot copy of the clustered Data ONTAP volume by entering the following command:

```
snapmirror show -destination-path vserver_name:volume_name
```

The Newest Snapshot field provides information about the latest Snapshot copy.

Example

In the following example, vs1(2014307946)_vol1.9 is the latest Snapshot copy of the transitioned volume vol1:

```
cluster1::> snapmirror show -destination-path vsv1:vol1
      Source Path: 10.225.243.7:vol_sat
      Destination Path: vs1:vol1
      Relationship Type: TDP
      SnapMirror Schedule: -
      Tries Limit: -
      Throttle (KB/sec): unlimited
      Mirror State: Snapmirrored
      Relationship Status: Transferring
      Transfer Snapshot: -
      Snapshot Progress: -
      Total Progress: 84KB
      Snapshot Checkpoint: -
      Newest Snapshot: vs1(2014307946)_vol1.9
      .....
```

2. Verify what the latest Snapshot copy is in the 7-Mode system.
 - a. Log in to the 7-Mode system.

- b. Verify whether the latest Snapshot copy on the Vserver also exists in the 7-Mode system by entering the following command:

```
snap list -V volume_name
```

Example

The following example shows that the latest Snapshot copy vs1(2014307946)_vol1.9 for Vserver vs1 does not exist on the 7-Mode volume:

```
system1> snap list vol1
Volume vol1
working...
```

%/used	%/total	date	name
0% (0%)	0% (0%)	Feb 28 20:00	hourly.0
0% (0%)	0% (0%)	Feb 28 16:00	hourly.1
0% (0%)	0% (0%)	Feb 28 12:00	hourly.2
0% (0%)	0% (0%)	Feb 28 08:00	hourly.3
0% (0%)	0% (0%)	Feb 27 00:00	nightly.0
0% (0%)	0% (0%)	Feb 26 20:00	hourly.4
0% (0%)	0% (0%)	Feb 26 16:00	hourly.5
0% (0%)	0% (0%)	Feb 25 03:56	vol1_snap1

3. If the Snapshot copy does not exist in the 7-Mode system, break the SnapMirror relationship from the cluster by entering the following command:

```
snapmirror break -destination-path vsserver_name:volume_name
```

Example

```
cluster1::>snapmirror break -destination-path vs1:vol1
```

4. Resynchronize the SnapMirror relationship by entering the following command:

```
snapmirror resync -destination-path vsserver_name:volume_name
```

Example

```
cluster1::>snapmirror resync -destination-path vs1:vol1
```

5. From the 7-Mode system, verify whether the latest Snapshot copy now exists in the 7-Mode system by entering the following command:

```
snap list -V volume_name
```

6. From the 7-Mode Transition Tool, run the transition complete operation again.

- a. Log in to the 7-Mode Transition Tool.
- b. Run the transition complete operation by entering the following command:

```
transition complete -s session_name [-m]
```

Transition pause operation failed

Details The transition pause operation fails.

Cause 1 The pause operation was run on a volume with a failed baseline transfer.

Corrective action 1 1. On the cluster, use the `snapmirror show` command to verify whether the baseline transfer failed for the volume.

The following example shows a failed baseline transfer for `vol1`, where the `Mirror State` is `empty`, `Last Transfer Type` is “initialize”, and `Last Transfer Error` provides the reason of failure for the baseline transfer:

```
cluster1::> snapmirror show -destination-path vs1:vol1
      Source Path: 192.0.55.28:vol1
      Destination Path: vs1:vol1_1
      Relationship Type: TDP
      SnapMirror Schedule: daily
      Tries Limit: -
      Throttle (KB/sec): unlimited
      Mirror State: -
      Relationship Status: Idle
      Transfer Snapshot: -
      Snapshot Progress: -
      Total Progress: -
      Snapshot Checkpoint: -
      Newest Snapshot: -
      Newest Snapshot Timestamp: -
      Exported Snapshot: -
      Exported Snapshot Timestamp: -
      Healthy: false
      Unhealthy Reason: Transfer failed.
      Constituent Relationship: false
      Destination Volume Node: cluster1-01
      Relationship ID: 5a175c4b-86ee-11e2-b061-123478563412
      Transfer Type: -
      Transfer Error: -
-
      Current Throttle: -
      Current Transfer Priority: -
      Last Transfer Type: initialize
not connect to the source host 192.0.55.28.
      Last Transfer Size: -
      Last Transfer Duration: -
      Last Transfer From: 192.0.55.28:vol1
      Last Transfer End Timestamp: 03/07 06:16:28
      Progress Last Updated: -
      Relationship Capability: 8.2 and above
      Lag Time: -
      SnapMirror Policy: DPDefault
```

Last Transfer Error: Could

2. Fix the cause for the last transfer error.
3. From the 7-Mode Transition Tool, start the baseline transfer by using the `transition start` command.
The baseline transfer is started.
4. Run the `transition pause` command after the baseline transfer has started and a restart checkpoint is created.

Cause 2 The pause operation is run immediately after running the transition start operation. To pause a session that is in the copy-baseline state, all the volumes in the session must have a restart checkpoint available, which is created some time after the transfer is started.

Corrective action 2

1. From the destination cluster, use the `snapmirror show` command with the `-snapshot-checkpoint` parameter to verify that a restart checkpoint exists for the volume.

```
cluster2::> snapmirror show -destination-path
dest_vserver:vol3 -fields snapshot-checkpoint
source-path          destination-path snapshot-checkpoint
-----
src_system:vol3      dest_vserver:vol3 50MB
```

2. From the 7-Mode Transition Tool, run the `transition pause` command.

Transition DR relationship not established if relationship is established using vFiler context

When transitioning the secondary volume of a volume SnapMirror relationship, if the host name of the primary 7-Mode system is not resolvable from the default vFiler context (vfiler0) of the secondary 7-Mode system, the transition DR relationship between the 7-Mode primary volume and clustered Data ONTAP secondary is not established.

Workaround

Edit the `/etc/hosts` file of the default vFiler unit of the secondary 7-Mode system and add the host entry for the primary 7-Mode system to the file.

This ensures that the host name of the primary 7-Mode controller resolves to the correct primary 7-Mode controller IP address from the default vFiler unit context. This is required even when the SnapMirror relationships are established from the non-default vFiler unit context of the secondary 7-Mode controller.

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

The tool displays the Another instance of NetApp 7-Mode Transition Tool Setup or Uninstall is running. Only one instance can run at a time. error message when using the 7-Mode Transition Tool installer or the uninstaller to perform any operations (including install, repair, reinstall, upgrade, uninstall).

Workaround

1. Open **Registry** by clicking **Start**.
2. In the **Search programs and files**, enter

RegEdit

and press Enter.

3. Once Registry Editor is open, locate the HKEY_LOCAL_MACHINE registry hive.
4. Expand the HKEY_LOCAL_MACHINE hive by clicking the expand icon to the left of the folder icon.
5. Continue to expand the registry keys and sub-keys until you reach the HKEY_LOCAL_MACHINE\SOFTWARE \NetApp \7-Mode Transition Tool.
6. Click the **7-Mode Transition Tool**.
7. Right-click INSTALLATION_STATUS and select **Modify...**
8. Change the value from INSTALLER_RUNNING to INSTALLED, and click **OK**.

Now you should be able to perform any operation using the NetApp 7-Mode Transition Tool installer or the uninstaller.

Note: Before performing the workaround, ensure that there are no open instances or dialog boxes related to the NetApp 7-Mode Transition Tool. Otherwise you may need to perform the steps again.

Error codes and corrective tasks for transition

Certain configurations are not transitioned by the 7-Mode Transition Tool because either these configurations are not supported or the tool fails to apply the configuration. In such cases, you should perform certain corrective tasks on the 7-Mode system or the cluster.

Tip: The precheck report provides a corrective action that you can perform after the transition, for all errors and warnings. You can view the job ID of the precheck operation to view the precheck results.

You can filter the issues, by selecting the options under **Apply Type Filter** and **Apply Tag Filter**.

- 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 display all messages related to security, select **Error**, **Warning**, **Informational**, and **Security Only**.

The following tables list the information source for completing the corrective tasks on the 7-Mode system and the cluster for each of the following feature areas:

- [CIFS](#) on page 167
- [NFS](#) on page 169
- [Name services](#) on page 171
- [Networking](#) on page 172
- [Volumes and Vservers](#) on page 173
- [Data protection](#) on page 174

CIFS

Error code	Corrective task	More information
20042	Configuring CIFS server on the 7-Mode storage system to use Active Directory style authentication	<i>Data ONTAP File Access and Protocols Management Guide for 7-Mode</i>
20062	Removing Storage-Level Access Guard from a 7-Mode volume before transition	• Man page of the <code>fsecurity remove-guard</code> command • <i>Data ONTAP File Access and Protocols Management Guide for 7-Mode</i>
20072	Removing Storage-Level Access Guard from a 7-Mode qtree before transition	

Error code	Corrective task	More information
20082	Removing character mapping on the 7-Mode volume before transition	- Man page of the charmap command <i>Data ONTAP File Access and Protocols Management Guide for 7-Mode</i>
20102	Editing the <code>/etc/usermap.cfg</code> file to remove the dot notation	<i>Data ONTAP File Access and Protocols Management Guide for 7-Mode</i>
20122	Adding CIFS license to the cluster	- Man page of the system license add command <i>Clustered Data ONTAP System Administration Guide for Cluster Administrators</i>
20132 20133	Configuring CIFS server on the Vserver	<i>Clustered Data ONTAP File Access Management Guide for CIFS</i>
20142 20042	Reconfiguring the CIFS server to use the Active Directory domain of the 7-Mode system	- Man page of the cifs setup command <i>Data ONTAP File Access and Protocols Management Guide for 7-Mode</i>
20213 20221	Creating an export rule for the export policy after completing the transition to allow CIFS access on the clustered Data ONTAP volume	- Man page of the vservers export-policy rule create command <i>Clustered Data ONTAP File Access Management Guide for CIFS</i>
20303 20304	Creating CIFS shares and applying ACLs to these shares	<i>Clustered Data ONTAP File Access Management Guide for CIFS</i>
20312	Configuring CIFS share-level ACLs for users or groups	
20412	Creating CIFS shares	

Error code	Corrective task	More information
20422 20423	Adding CIFS to the allowed list of protocols on the Vserver	- Man page of the <code>vserver modify</code> command <i>Clustered Data ONTAP File Access Management Guide for CIFS</i>
20432	Modifying the administrative status of the CIFS server on the Vserver	- Man page of the <code>vserver modify</code> command <i>Clustered Data ONTAP File Access Management Guide for CIFS</i>
20442	Configuring FPolicy on the Vserver	<i>Clustered Data ONTAP File Access Management Guide for CIFS</i>
20512	Applying group policy objects (GPOs) to the CIFS server	
20702	Configuring CIFS home-directory paths	
20772	Creating local UNIX users and groups	
20602	Configuring Branch Cache	- Man page of the <code>vserver cifs branchcache create</code> command <i>Clustered Data ONTAP File Access Management Guide for CIFS</i>
20612	Enabling CIFS over IPv6	<i>Clustered Data ONTAP File Access Management Guide for CIFS</i> <i>Clustered Data ONTAP Network Management Guide</i>
20622	Configuring CIFS auditing on the Vserver	<i>Clustered Data ONTAP File Access Management Guide for CIFS</i>
20632	Enabling SMB	

NFS

Error code	Corrective task	More information
10012	Enabling NFS license on the cluster	Man page of the <code>system license add</code> command

Error code	Corrective task	More information
10031 10041 10072 10082 10092 10102 10112 10122 10132 10052	Editing the <code>/etc/exports</code> file to match the entries supported in clustered Data ONTAP	• <i>Data ONTAP File Access and Protocols Management Guide for 7-Mode</i> • <i>Clustered Data ONTAP File Access Management Guide for NFS</i> • <i>Clustered Data ONTAP Commands: Manual Page Reference</i>
10201	Configuring NFS over IPv6	<i>Clustered Data ONTAP File Access Management Guide for NFS</i> <i>Clustered Data ONTAP Network Management Guide</i>
10221 10223 10252 10272	Configuring an Active Directory-based Kerberos server	<i>Clustered Data ONTAP File Access Management Guide for NFS</i>
10222 10224 10262	Configuring an UNIX-based Kerberos server	
10227 10372	Creating a Kerberos realm configuration	
10241 10282	Configuring Kerberos on the LIFs	
		• Man page of the <code>vserver nfs kerberos-config modify</code> command • <i>Clustered Data ONTAP File Access Management Guide for NFS</i>

Error code	Corrective task	More information
10393 10471 10472 10474	Creating an export rule for the default export policy on the Vserver	- Man page of the vserver export-policy rule create command <i>Clustered Data ONTAP File Access Management Guide for NFS</i>
10752	Enabling NFS auditing on the Vserver	<i>Clustered Data ONTAP File Access Management Guide for NFS</i>
10212	Adding NFS to the allowed list of protocols on the Vserver	- Man page of the vserver modify command <i>Clustered Data ONTAP File Access Management Guide for NFS</i>

Name services

Error code	Corrective task	More information
99172	Configuring protocols and services	<i>Clustered Data ONTAP File Access Management Guide for NFS</i> <i>Clustered Data ONTAP File Access Management Guide for CIFS</i> <i>Clustered Data ONTAP Network Management Guide</i>
40105 40107	Configuring DNS domains after transition	<i>Clustered Data ONTAP Network Management Guide</i>

Error code	Corrective task	More information
40143 40144	Configuring NIS after transition	<i>Clustered Data ONTAP File Access Management Guide for NFS</i>
40151	Configuring LDAP on the cluster	
40162 40172	Modifying an LDAP schema	
40182 40183	Creating an LDAP client configuration	
40063	Configuring NIS domain on the Vserver with specific NIS servers	
40202	Applying LDAP client configuration to the Vserver	
40191 40591 40592	Configuring LDAP bind password	Man page of the vserver services ldap client modify command
40312	Configuring netgroups on the Vserver	• Man page of the vserver services netgroup load command • <i>Clustered Data ONTAP File Access Management Guide for NFS</i>

Networking

Error code	Corrective task	More information
10233 10234 96702 96752 99162 30092	Configuring data LIFs and IP addresses on the Vserver after transition	<i>Clustered Data ONTAP Network Management Guide</i>
30021	Configuring interface groups on the cluster	• Man page of the network port ifgrp create command • <i>Clustered Data ONTAP Network Management Guide</i>

Error code	Corrective task	More information
30031	Configuring VLANs on the cluster	- Man page of the network port <code>vlan create</code> command <i>Clustered Data ONTAP Network Management Guide</i>
30071	Configuring interface groups or VLANs on the cluster	- Man page of the network port <code>ifgrp create</code> command - Man page of the network port <code>vlan create</code> command <i>Clustered Data ONTAP Network Management Guide</i>
30041	Creating routing groups and routes	<i>Clustered Data ONTAP Network Management Guide</i>
30512	Configuring IPv6 on the cluster	

Volumes and Vservers

Error code	Corrective task	More information
10321 10332 10341	Creating junction paths for the volumes in the Vserver	<i>Clustered Data ONTAP Logical Storage Management Guide</i>
95312 95372	Adding disks to the aggregate	- Man page of the storage aggregate <code>add-disks</code> command <i>Clustered Data ONTAP Logical Storage Management Guide</i>
60323	Configuring deduplication and compression on the volumes after transition	<i>Clustered Data ONTAP Logical Storage Management Guide</i>
95432	Increasing the maximum number of volumes that can be created for a Vserver	Man page of the vserver <code>modify</code> command

Error code	Corrective task	More information
70102 70112 70122 70123 70182 70192 90242	Enabling quotas and applying quota rules to volumes after transition	<i>Clustered Data ONTAP Logical Storage Management Guide</i>
96803	Taking the volume offline on the 7-Mode system	Man page of the vol offline command

Data protection

Item	Corrective task	More information
70132 70202	Creating a Snapshot copy schedule	<i>Clustered Data ONTAP Data Protection Guide</i>
70212	Applying space reservation configuration to the clustered Data ONTAP volumes	<i>Clustered Data ONTAP Data Protection Guide</i>
90252	Deleting the Snapshot copies from the 7-Mode system	- Man page of the snap delete command <i>Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode</i>
90303	Breaking the qtree SnapVault relationships on the 7-Mode system	<i>Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode</i>
70073	Breaking the qtree SnapMirror relationship on the 7-Mode destination system	<i>Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode</i>
91151 91161	Configuring SnapMirror schedules on the destination after transition	<i>Clustered Data ONTAP Data Protection Guide</i>
60352	Configuring NDMP backup policies	<i>Clustered Data ONTAP Data Protection Tape Backup and Recovery Guide</i>

Item	Corrective task	More information
90843	1. Breaking the SnapMirror relationship manually from the destination cluster 2. Running the transition complete operation again	1. Man page of the snapmirror break command 2. Completing a transition subproject on page 76
90853	Running the transition complete operation again	Completing a transition subproject on page 76

Related information

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

Impact of Data ONTAP features on transition

You should be aware of the impact of some of the Data ONTAP features, such as tape backup and high-availability configurations, during transition.

Impact of takeover and giveback on transition

Transition operations, such as transition prepare, start, pause, resume, or complete, fail during a controller takeover or giveback.

If a transition operation fails due to a takeover, you must wait for the giveback to complete and run the transition operation again.

If a controller takeover occurs during a baseline transfer, the transfer fails. To resume the baseline transfer from the point where it was aborted, you must wait for the giveback to complete and then run the transition start operation again.

About backing up volumes in transition to tape

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

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

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

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

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

- When a SnapMirror relationship is established between the clustered Data ONTAP primary and secondary volumes

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

Glossary

A

ACL Access control list.

AD Active Directory.

administration host A client computer that is used to manage a storage system through a Telnet or Remote Shell connection.

aggregate A grouping of physical storage resources (disks or array LUNs) that provides storage to volumes associated with the aggregate. Aggregates provide the ability to control the RAID configuration for all associated volumes.

C

CFO See *controller failover (CFO)* or *cluster failover (CFO)*.

CIFS share

- In Data ONTAP, a directory or directory structure that has been made available to network users and can be mapped to a drive letter on a CIFS client. Also known simply as a *share*.
- In OnCommand Insight (formerly SANscreen suite), a service exposed from a NAS device to provide file-based storage through the CIFS protocol. CIFS is mostly used for Microsoft Windows clients, but many other operating systems can access CIFS shares as well.

CLI command-line interface. The storage system prompt is an example of a command-line interface.

client A workstation or PC in a client-server architecture; that is, a computer system or process that requests services from and accepts the responses of another computer system or process.

cluster

- In clustered Data ONTAP 8.x, a group of connected nodes (storage systems) that share a namespace and that you can manage as a single virtual server or multiple virtual servers, providing performance, reliability, and scalability benefits.
- In the Data ONTAP 7.1 release family and earlier releases, a pair of storage systems (sometimes called *nodes*) configured to serve data for each other if one of the two systems stops functioning.
- In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.
- For some storage array vendors, *cluster* refers to the hardware component on which host adapters and ports are located. Some storage array vendors refer to this component as a *controller*.

cluster failover (CFO)	In Data ONTAP 7.1.x and earlier, the method of ensuring data availability by transferring the data service of a failed node to another node in an HA pair. Transfer of data service is often transparent to users and applications. In Data ONTAP 7.2 and later, and in Data ONTAP operating in 7-Mode, the failover method is called <i>controller failover</i> . In clustered Data ONTAP, the failover method is called <i>storage failover</i> .
cluster Vserver	Former name for a data Vserver; see data Vserver.
Common Internet File System (CIFS)	Microsoft's file-sharing networking protocol that evolved from SMB.
D	
data protection	The combination of processes and systems that ensures that data is available, accessible, secure, recoverable, and in compliance with applicable requirements. Can imply elements of availability, backup and recovery, disaster recovery, security, and compliance.
data Vserver	Formerly known as cluster Vserver. In clustered Data ONTAP, a Vserver that facilitates data access from the cluster; the hardware and storage resources of the cluster are dynamically shared by data Vservers within a cluster.
dedupe	See <i>deduplication</i> .
deduplication	The consolidation of blocks of duplicate data into single blocks to store more information using less storage space.
destination	The volume or qtree to which data is mirrored in a SnapMirror relationship. The destination volume or qtree must be the same as the source volume or qtree.
disaster recovery (DR)	The recovery of data, access to data, and associated comprehensive processing of setting up the equipment and work space of a redundant site to continue business operations after a loss of use of all or part of a data center. This involves not only recovering an essential set of data but also an essential set of all the hardware and software to continue the processing of that data and business. Any disaster recovery might involve some amount of down time.
DNS	See <i>Domain Name System (DNS)</i> .
Domain Controller (DC)	The computer running NT Server that is contacted when a request comes in to log in a CIFS client or validate a request from an NFS client accessing a file with NTFS-style security. A given domain has one Primary Domain Controller (PDC) and optionally one or more Backup Domain Controllers (BDCs).
Domain Name System (DNS)	An Internet service for finding computers on an IP network.

DR	See <i>disaster recovery (DR)</i> .
E	
export	The NFS mechanism used to define which areas of the storage system are to be made available to NFS clients.
F	
failover	A process that reroutes data immediately to an alternate path so that services remain uninterrupted, if a physical disruption to a network component occurs. Failover applies to clustering and to multiple paths to storage. In the case of clustering, one or more services, such as the Exchange service, is moved to a standby server. In the case of multiple paths to storage, a path failure results in data being rerouted to a different physical connection to the storage.
FlexClone	A NetApp software technology that allows instant replication of data files, LUNs, and volumes without requiring additional storage space. Each cloned file, LUN, or volume is a virtual copy of the original data.
FPolicy	The Data ONTAP proprietary file policy feature that provides the ability to control access permissions based on file properties, such as file type.
FQDN	See <i>Fully Qualified Domain Name (FQDN)</i> .
Fully Qualified Domain Name (FQDN)	The complete name of a specific computer on the Internet, consisting of the computer's host name and its domain name.
G	
GID	See <i>Group ID (GID)</i> .
giveback	The technology that enables two storage systems to return control of each other's data after the issues that caused a controller failover are resolved.
global namespace	See <i>namespace</i> .
group	In Data ONTAP operating in 7-Mode, a group of users defined in the storage system's <code>/etc/group</code> file.
Group ID (GID)	The number used by UNIX systems to identify groups.
H	
HA (high availability)	- In Data ONTAP 8.x, the recovery capability provided by a pair of nodes (storage systems), called an <i>HA pair</i>, that are configured to serve data for each other if one of the two nodes stops functioning. - In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an <i>active/active configuration</i>.

HA pair	- In Data ONTAP 8.x, a pair of nodes whose controllers are configured to serve data for each other if one of the two nodes stops functioning. Depending on the system model, both controllers can be in a single chassis, or one controller can be in one chassis and the other controller can be in a separate chassis. - In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an <i>active/active configuration</i>.
L	
LIF	logical interface. Formerly known as <i>VIF</i> (virtual interface) in Data ONTAP GX. A logical network interface, representing a network access point to a node. LIFs currently correspond to IP addresses, but could be implemented by any interconnect. A LIF is generally bound to a physical network port; that is, an Ethernet port. LIFs can fail over to other physical ports (potentially on other nodes) based on policies interpreted by the LIF manager.
local group	A collection of user accounts and global groups in a Windows NT domain. Unlike <i>global groups</i> , local groups can contain members from other, trusted domains.
LUN (Logical Unit Number)	The identifier of an FC or iSCSI logical unit. A logical unit typically corresponds to a storage volume and is represented within a computer operating system as a device.
M	
MultiStore	In Data ONTAP operating in 7-Mode, an optional software product that enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network.
migration	The process of moving a set of data from one storage location to another storage location. See also <i>offline migration</i> and <i>online migration</i> .
N	
namespace	In network-attached storage (NAS) environments, a collection of files and path names to the files.
NAS	network-attached storage.
NDMP	Network Data Management Protocol. A protocol that allows storage systems to communicate with backup applications and provides capabilities for controlling the robotics of multiple tape backup devices.
Network File System (NFS) export	A service exposed from a NAS device to provide file-based storage through the NFS protocol. NFS is mostly used for UNIX-like operating systems, but other operating systems can access NFS exports as well.
NFS	Network File System.

normal mode	The state of a storage system when there is no takeover in the HA configuration.
O	
offline migration	A data migration process in which the data being migrated is not available to users during the migration cutover operation. All operations in the migration process must be initiated manually.
online migration	A data migration process in which the data being migrated is available to users at all times during the migration. The online migration process allows you to enable the automatic start of the migration cutover operation when the migration start operation is complete.
P	
partner mode	The method you use to communicate through the command-line interface with a virtual storage system during a takeover.
partner node	From the point of view of the local node (storage system), the other node in a HA configuration.
(PC)NFS	Personal Computer Network File System.
pcnfsd	A storage system daemon that permits PCs to mount storage system file systems. The corresponding PC client software is called <i>PC-NFS</i> .
Q	
qtree	A special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes.
R	
RIP	Routing Information Protocol.
S	
SAN	storage area network.
share	A directory or directory structure that has been made available to network users and can be mapped to a drive letter on a CIFS client. Also known as a <i>CIFS share</i> .
SFO	See <i>storage failover (SFO)</i> .
SID	Security identifier used by the Windows operating system.
Snapshot copy	An online, read-only copy of an entire file system that protects against accidental deletions or modifications of files without duplicating file contents. Snapshot copies enable users to restore files and to back up the storage system to tape while the storage system is in use.
SSH	Secure Shell.

SSL	Secure Sockets Layer.
storage controller	The component of a storage system that runs the operating system and controls its disk subsystem. Storage controllers are also sometimes called <i>controllers</i> , <i>storage appliances</i> , <i>appliances</i> , <i>storage engines</i> , <i>heads</i> , <i>CPU modules</i> , or <i>controller modules</i> .
storage efficiency	The ratio of usable capacity to effective used capacity, accounting for efficiency returns. Calculated by the effective used capacity / the usable capacity.
storage failover (SFO)	In clustered Data ONTAP, the method of ensuring data availability by transferring the data service of a failed node to another node in an HA pair. Transfer of data service is often transparent to users and applications. In Data ONTAP 7.2 and later, and in Data ONTAP operating in 7-Mode, the failover method is called <i>controller failover</i> .
storage system	The hardware device running Data ONTAP that receives data from and sends data to native disk shelves, third-party storage, or both. Storage systems include a controller component and an internal or external disk storage subsystem component. Storage systems are sometimes referred to as <i>filers</i> , <i>appliances</i> , <i>storage appliances</i> , <i>V-Series systems</i> , or <i>systems</i> .
T	
takeover	The emulation of the failed node identity by the takeover node in an HA pair; the opposite of <i>giveback</i> .
takeover mode	The method you use to interact with a node (storage system) when it has taken over its partner. The console prompt indicates when the node is in takeover mode.
takeover node	A node (storage system) that remains in operation after the other node stops working and that hosts a virtual node that manages access to the failed node disk shelves and network connections. The takeover node maintains its own identity and the virtual node maintains the failed node identity.
U	
UID	user identification number.
Unicode	A 16-bit character set standard. It was designed and is maintained by the nonprofit consortium Unicode Inc.
V	
vFiler unit	In Data ONTAP operating in 7-Mode, a virtual storage system that you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network.
VLAN	virtual local area network.

volume	A file system.
Vserver	(Known as “Storage Virtual Machine (SVM)” in clustered Data ONTAP 8.2.1 and later.) A virtual machine that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of Vservers— <i>admin</i> , <i>node</i> , and <i>cluster</i> (“cluster Vserver” is called “data Vserver” in Data ONTAP 8.2)—but unless there is a specific need to identify the type of Vserver, “Vserver” usually refers to the cluster/data Vserver.
W	
WAFL	Write Anywhere File Layout. A file system designed for the storage system to optimize write performance.
WAFL External Cache	On a storage system that has a Performance Acceleration Module (PAM), Flash Cache, or Flash Cache 2 module installed, this cache improves storage system performance by reducing the number of disk reads. Sometimes referred to as <i>WAFL extended cache</i> .
WINS	Windows Internet Name Service.
workgroup	A collection of computers running Windows NT or Windows for Workgroups that is grouped for browsing and sharing.

Copyright information

Copyright © 1994–2014 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

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

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Customer Fitness, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

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

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

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

Index

7-Mode

- data migration to clustered Data ONTAP [7](#)
- preparation for transition [18](#)
- transitioning volumes and configuration to clustered Data ONTAP [7](#)

7-Mode Transition Tool

- CLI [32](#)
- configuration file [49](#)
- error messages [147](#)
- GUI [32](#)
- installing [27](#), [28](#)
- log files [146](#)
- logging in using the web interface [69](#)
- network port requirements [19](#)
- system requirements [27](#)
- uninstalling [27](#), [30](#)
- upgrading [27](#)
- web interface
 - logging in to 7-Mode Transition Tool [69](#)
 - what it does [7](#)

A

- aborting
 - a transition session [141](#)
- access is denied [149](#)
- audit log path of the Vserver
 - setting [138](#)

B

- backup of volumes in transition
 - using dump [176](#)
- best practices [14](#)
- bidirectional SnapMirror relationships
 - how to transition [85](#)
- browsers
 - supported versions for the web interface [27](#)

C

- CIFS
 - prerequisites to transition [20](#)
 - scenario for transition [85](#)
- client access

- volumes in data protection relationship [36](#)
- client access phase
 - transition of stand-alone volumes [33](#)
- clustered Data ONTAP
 - transitioning data from 7-Mode [7](#)
- clusters
 - adding [70](#)
 - preparing for transition [15](#)
- commands to manage
 - credentials [143](#)
 - jobs [143](#)
 - LIFs [143](#)
 - SnapMirror relationships [143](#)
 - transition operations [143](#)
 - transition sessions [143](#)
 - volume pairs [143](#)
- compatibility
 - verifying, for transition [135](#)
- completing projects
 - prerequisites [68](#)
- configuration file
 - described [49](#)
- considerations for CIFS [55](#)
- considerations for NFS
 - considerations for NFS transition
 - considerations for NFS exports transition [51](#)
 - considerations for NFS exports transition [51](#)
 - considerations for NFS transition [51](#)
- controllers
 - adding [70](#)
 - assessing [71](#)
- copy-unknown error [159](#)
- creating
 - a transition session [131](#)
- creating projects
 - prerequisites [68](#)
- credentials
 - commands to manage [143](#)
- cutover
 - completing transition for [139](#)
 - guidelines for performing [47](#)
- cutover phase
 - transition of stand-alone volumes [33](#)
 - volumes in a data protection relationship [36](#)

D

- dashboard
 - aborting subprojects [81](#)
 - prechecks [77](#)
- data copy
 - initiating [136](#)
- data copy phase
 - transition of stand-alone volumes [33](#)
 - volumes in a data protection relationship [36](#)
- Data ONTAP
 - 7-Mode versions supported [12](#)
- deleting
 - a Snapshot copy [161](#)
 - a transition session [141](#)
- displaying
 - session details [140](#)
- dump support for transition
 - understanding [176](#)

E

- error codes
 - corrective tasks for [167](#)
- error message
 - 7-Mode Transition Tool installer is blocked
 - installation or uninstallation of the 7-Mode Transition Tool is blocked [165](#)
 - cannot add credentials to controllers
 - cannot add credentials to clusters [83](#)
- error messages
 - 7-Mode Transition Tool [147](#)
 - communication issues
 - host not reachable [152](#), [154](#)
 - host timed out [152](#), [154](#)
 - error occurred with Cluster-Mode storage system [150](#)
 - error onlining volume [150](#)
 - exception in thread "main"
 - java.lang.OutOfMemoryError [158](#)
 - failed to create specified directory [147](#)
 - failed to fetch volume information [150](#)
 - final update failed [161](#)
 - host not found in credential bank [149](#)
 - network connectivity issues with 7-Mode system [152](#)
 - network connectivity issues with cluster [154](#)
 - no content to map to Object due to end of input [158](#)
 - resolving BUILIN CIFS name conflict [148](#)
 - there is not enough space on the disk [158](#)
 - unexpected error occurred while detecting java [158](#)

when opening the CLI [149](#)

F

- failed baseline transfer [164](#)
- features
 - not supported for transition [8](#)
- FileNotFoundException exception [149](#)

G

- giveback
 - impact on transition [176](#)

I

- installing
 - 7-Mode Transition Tool [27](#), [28](#)

J

- java.io.FileNotFoundException [149](#)
- jobs
 - commands to manage [143](#)
 - viewing transition [141](#)

L

- licenses
 - requirements for transition [12](#)
- LIF
 - See* logical interfaces
- LIFs
 - commands to manage [143](#)
- limits
 - for transition [48](#)
- log files
 - 7-Mode Transition Tool [146](#)
- logical interfaces
 - managing [79](#)
 - See also* LIFs
- logs
 - collecting logs
 - tool logs [82](#)

N

- network connectivity
 - issues with the 7-Mode system [152](#)

- issues with the cluster [154](#)
- network ports
 - requirements [19](#)
- NFS
 - prerequisites to transition [20](#)
- NFS only transition
 - NFS only configuration transition [55](#)

P

- pause failed [164](#)
- pausing
 - a transition session [138](#)
- phases
 - stand-alone volumes, transition [33](#)
 - volumes in data protection relationship, transition [36](#)
- ports
 - requirements [19](#)
- prechecks
 - error [49](#)
 - information [49](#)
 - notification [49](#)
 - running [77](#), [135](#)
 - severity levels [49](#)
 - warning [49](#)
- precutover phase
 - transition of stand-alone volumes [33](#)
 - volumes in a data protection relationship [36](#)
- preparation
 - for transition, checklist [22](#)
- preparation phase
 - transition of stand-alone volumes [33](#)
 - volumes in a data protection relationship [36](#)
- prerequisites
 - creating projects [68](#)
 - running projects [68](#)
 - starting projects [68](#)
- primary volumes
 - transitioning [119](#)
- project creation
 - stand-alone [58](#)
 - understanding workflow for [58](#)
- projects
 - restrictions [66](#)
 - understanding workflows for creating [61](#)
 - understanding workflows for starting [61](#)
 - workflows for creating [58](#)

R

- removing
 - volume from session
 - removing
 - LIF from session [80](#)
- requirements
 - 7-Mode Transition Tool [19](#)
 - for the tool, system [27](#)
 - for transitioning NFS and CIFS configurations [20](#)
 - licenses for transition [12](#)
- restart checkpoint unavailable [164](#)
- restoring
 - a Snapshot copy [162](#)
- restrictions
 - on transition [14](#)
 - projects [66](#)
- resuming
 - a transition session [138](#)
- running projects
 - prerequisites [68](#)

S

- SAN
 - unsupported features [8](#)
- secondary subproject
 - aborting [81](#)
- secondary volumes
 - transitioning [107](#)
- session states
 - copy-unknown [159](#)
- sessions
 - primary [33](#)
 - scenarios for creating [85](#)
 - secondary [33](#)
 - stand-alone [33](#)
- SnapMirror relationships
 - commands to manage [143](#)
 - completing [64](#)
 - considerations for transition [45](#)
 - creating a project that transitions [60](#)
 - scenario for transition [85](#)
 - starting [64](#)
 - transitioning, volume [103](#)
 - workflow [60](#), [64](#)
 - worksheet to use for transitioning [103](#)
- SnapMirror transfers
 - pausing and resuming [138](#)
- Snapshot copies

- deleting [161](#)
 - restoring [162](#)
 - stand-alone subproject
 - aborting [81](#)
 - stand-alone volumes
 - transitioning [87](#)
 - stand-alone workflows
 - completing [63](#)
 - starting [63](#)
 - starting projects
 - prerequisites [68](#)
 - stopping
 - a transition session [141](#)
 - subprojects
 - aborting [81](#)
 - editing [78](#)
 - managing [77](#)
 - pausing and resuming [80](#)
 - primary [33](#)
 - secondary [33](#)
 - stand-alone [33](#)
 - starting [75](#)
 - system requirements
 - for the tool [27](#)
- ## T
- takeover
 - impact on transition [176](#)
 - transition
 - 7-Mode prerequisites [18](#)
 - blocking features [8](#)
 - checklist [22](#)
 - CIFS configuration prerequisites [20](#)
 - commands to manage [143](#)
 - concurrent transfers [45](#)
 - configuration file for [49](#)
 - considerations about entities [13](#)
 - considerations for local users and groups [55](#)
 - considerations for SnapMirror [45](#)
 - considerations for transition [51](#)
 - data copy over multiple paths [45](#)
 - Data ONTAP version requirements [12](#)
 - error codes and corrective tasks [167](#)
 - examples for creating sessions [85](#)
 - guidelines for cutover [47](#)
 - identifying the workloads [13](#)
 - impact of takeover and giveback [176](#)
 - jobs, viewing results of [141](#)
 - jobs, viewing status of [141](#)
 - license requirements [12](#)
 - limits [48](#)
 - local users and groups
 - considerations during CIFS transition [55](#)
 - NFS configuration prerequisites [20](#)
 - phases of [10, 36](#)
 - prechecks [49](#)
 - preparing for [12](#)
 - preparing the 7-Mode system [18](#)
 - preparing the cluster for [15](#)
 - project definition [32](#)
 - restrictions on [14](#)
 - session definition [32](#)
 - sessions [33](#)
 - SnapMirror multipath [45](#)
 - stand-alone volumes [33](#)
 - subproject definition [32](#)
 - subprojects [33](#)
 - supported 7-Mode versions [12](#)
 - terminology [32](#)
 - troubleshooting issues [146](#)
 - unsupported features [8](#)
 - using 7-Mode Transition Tool [7](#)
 - using the 7-Mode Transition Tool [32](#)
 - using the 7-Mode Transition Tool (CLI) [85](#)
 - volumes in a data protection relationship [36](#)
 - what you can transition [9](#)
 - transition complete
 - failure [161](#)
 - modes to use for [41](#)
 - normal mode [41, 139](#)
 - test mode [41, 139](#)
 - transition DR failed
 - relationship from vFiler context [165](#)
 - transition DR relationship not established [165](#)
 - transition operations
 - commands to manage [143](#)
 - transition pause
 - failure [164](#)
 - transition phases
 - cutover [10](#)
 - data copy [10](#)
 - preparation [10](#)
 - transition process
 - understanding [57](#)
 - transition project
 - completing [76](#)
 - creating [72](#)
 - deleting [82](#)
 - editing [78](#)

- managing [77](#)
- overview [57](#)
- requirements and restrictions [66](#)
- running [69](#)
- starting [75](#)
- workflow [69](#)
- transition projects
 - prechecks [77](#)
- transition sessions
 - aborting [141](#)
 - commands to manage [143](#)
 - completing [139](#)
 - creating [131](#)
 - deleting [141](#)
 - displaying details about [140](#)
 - initiating data copy [136](#)
 - managing [130](#)
 - pausing [138](#)
 - resuming [138](#)
 - starting [136](#)
 - states of a [42](#)
 - verifying compatibility [135](#)
- transition start failed [150](#)
- transition states
 - aborted [42](#)
 - complete [42](#)
 - copy baseline [42](#)
 - copy pause [42](#)
 - copy update [42](#)
 - copy-unknown [42](#)
 - cutover [42](#)
 - deleted [42](#)
 - preparation [42](#)
- transition subprojects
 - completing [76](#)
 - editing [78](#)
 - pausing [80](#)
 - requirements and restrictions [66](#)
 - resuming [80](#)
 - states of a [42](#)
- transition-tool.conf
 - described [49](#)
- transitioning
 - from 7-Mode to clustered Data ONTAP [7](#)
 - primary volumes [119](#)
 - secondary volumes [107](#)
 - SnapMirror relationship, setup worksheet [103](#)
 - stand-alone volume, setup worksheet [100](#)
 - stand-alone volumes [87](#)

- volume SnapMirror relationships [103](#)
- troubleshooting
 - final update failed [161](#), [162](#)
 - low memory [158](#)
 - out of memory [158](#)
 - transition issues [146](#)
- types
 - of transition [9](#)

U

- uninstalling
 - 7-Mode Transition Tool [30](#)
- unsupported features
 - for transition [8](#)
 - SAN [8](#)
- upgrading
 - 7-Mode transition tool [29](#)

V

- volume in inconsistent state [150](#)
- volume pairs
 - commands to manage [143](#)
- volume SnapMirror relationships
 - transitioning [103](#)
- volumes
 - adding [72](#)
 - LIF
 - configuring [72](#)
 - logical interface
 - configuring [72](#)
 - number of volumes to transition [48](#)
 - scenarios for selecting them for a session [85](#)
 - transitioning primary [119](#)
 - transitioning secondary [107](#)
 - transitioning stand-alone volumes [87](#)
 - types of transition [9](#)
- volumes in transition
 - backing up [176](#)

W

- web interface
 - Dashboard [61](#)
 - troubleshooting [83](#)
- worksheet
 - to use while transitioning stand-alone volumes [100](#)