



Updated for 8.2.2

Clustered Data ONTAP[®] 8.2

7-Mode Data Transition Using SnapMirror[®]



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

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

Part number: 215-09220_A0
July 2014

Contents

Transitioning 7-Mode volumes using SnapMirror	4
Planning for transition	4
Features and volumes not supported for transition	5
7-Mode version requirements for transition	6
Considerations for using SnapMirror for transition	6
Preparing for transition	7
License requirements for transition	7
Preparing the 7-Mode system for transition	7
Preparing the cluster for transition	9
Creating a transition peer relationship	10
Transitioning volumes	11
Transitioning a stand-alone volume	12
Transitioning a volume SnapMirror relationship	17
Resuming a failed SnapMirror baseline transfer	31
Recovering from a disaster at the 7-Mode site during transition	32
Redirecting clients to the clustered Data ONTAP secondary volume after a disaster	33
Transitioning a stand-alone volume manually	33
Redirecting clients to the clustered Data ONTAP primary volume	38
Copyright information	41
Trademark information	42
How to send your comments	43
Index	44

Transitioning 7-Mode volumes using SnapMirror

You can copy data from 7-Mode volumes to clustered Data ONTAP volumes by using clustered Data ONTAP SnapMirror commands. You must then set up the protocols, services, and other configuration on the cluster after transition is complete.

About this task

This procedure provides the high-level tasks you have to perform for transition using SnapMirror.

Attention: You can transition only volumes in a NAS (CIFS and NFS) environment to clustered Data ONTAP.

Steps

1. *Verify that the volumes you plan to transition are supported for transition.*
2. *Prepare the 7-Mode system for transition.*
3. *Prepare the cluster for transition.*
4. *Create a transition peer relationship between the 7-Mode system as the source and the SVM as the destination.*
5. *Copy the data from the 7-Mode volume to the clustered Data ONTAP volume by creating a SnapMirror relationship between the two volumes.*

After you finish

After data migration finishes, you must perform the following tasks:

- Create a data LIF on the SVM to enable client access.
For details about this task, see the *Clustered Data ONTAP Network Management Guide*.
- Configure protocols, networking, and services on the SVM.
For details about these tasks, see the *Clustered Data ONTAP Network Management Guide*, *Clustered Data ONTAP File Access and Protocols Management Guide*, and *Clustered Data ONTAP File Access Management Guide for CIFS*.

Related information

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

Planning for transition

Before copying data from 7-Mode volumes to clustered Data ONTAP volumes, you must understand when to use SnapMirror to perform the transition and review the information for supported 7-Mode

versions and supported volumes for transition. You must also be aware of certain transition considerations.

Note: You must review the *Clustered Data ONTAP Release Notes* for the transition target release for a listing of the transition issues.

You can use SnapMirror to transition data in the following scenarios:

- 7-Mode Transition Tool does not support your requirements for transition; for example, 7-Mode Transition Tool requires a Windows host that might be unavailable in your environment.
7-Mode Transition Tool is the recommended way to transition 7-Mode volumes because the tool provides prechecks to verify the feasibility of transition and migrates all protocol, network, and services configurations along with the data.
You can install and use 7-Mode Transition Tool to perform the prechecks for transitioning and then use SnapMirror commands to perform the data migration from the 7-Mode volume to the clustered Data ONTAP volume.
- The cluster and Storage Virtual Machine (SVM) are already configured and only the data has to be transitioned from the 7-Mode volumes to the clustered Data ONTAP volumes.

Features and volumes not supported for transition

You cannot transition certain 7-Mode volumes, such as traditional volumes or SnapLock volumes and certain 7-Mode features, such as synchronous SnapMirror relationships, because some features might not be available in clustered Data ONTAP.

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

You cannot transition the following 7-Mode volumes or configurations:

- Volumes with LUNs and LUN clones (SAN transition)
- Restricted or offline volumes
- Traditional volumes
- Volumes with NFS-to-CIFS character mapping (charmap)
- Volumes with Storage-Level Access Guard configurations
- SnapLock volumes
- FlexCache volumes
- 7-Mode volumes in a 64-bit aggregate to a 32-bit aggregate in clustered Data ONTAP
- FlexClone volumes
FlexClone volumes can be transitioned as FlexVol volumes, but the 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
- Synchronous SnapMirror configuration
- Qtree SnapMirror relationships
- IPv6 configurations

- SnapVault relationships

7-Mode 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 systems.

You can transition volumes from the following 7-Mode versions to clustered Data ONTAP systems:

- 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
- Data ONTAP 8.1.2
- Data ONTAP 8.1.3
- Data ONTAP 8.2

Considerations for using SnapMirror for transition

You must be aware of certain considerations when running transition operations simultaneously with SnapMirror or SnapVault operations occurring in the 7-Mode system, such as the maximum number of concurrent SnapMirror transfers, data copy schedules, and using multiple paths for transition.

Maximum number of concurrent SnapMirror transfers

During transition, the maximum number of concurrent SnapMirror transfers supported on the 7-Mode and clustered Data ONTAP systems depends on the number of volume SnapMirror replication operations allowed for a specific storage system model.

For information about the maximum number of concurrent volume SnapMirror transfers for your system model, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

Data copy schedules

Data copy schedules for transition operations should not overlap with existing schedules for SnapMirror or SnapVault operations running on the 7-Mode system.

Using multiple paths for transition

You can specify two paths for transition by using a data copy IP address and a multipath IP address. However, both paths can be used only for load-balancing, not for failover.

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

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

Related information

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

Preparing for transition

Before you start the transition, you must prepare the 7-Mode storage system and cluster before transitioning 7-Mode volumes to clustered Data ONTAP. You must also create a transition peer relationship between the 7-Mode system and the Storage Virtual Machine (SVM).

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. If you are transitioning a 7-Mode volume SnapMirror relationship, SnapMirror licenses are also required on the source and destination clusters.

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

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

Related information

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

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 and the 7-Mode system to communicate with the target cluster.

Before you begin

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

8 | 7-Mode Data Transition Using SnapMirror

About this task

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

Steps

1. Add and enable the SnapMirror license on the 7-Mode system.
 - a. Add the SnapMirror license on the 7-Mode system by using the following command:

```
license add xxxxxxxx
```

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

```
options snapmirror.enable on
```
2. Configure the 7-Mode system and the target cluster to communicate with each other by choosing one of the following options:
 - Set the `snapmirror.access` option to `all`.
 - Set the value of the `snapmirror.access` option to the IP addresses of all the intercluster LIFs on the cluster.
 - If the `snapmirror.access` option is `legacy` and the `snapmirror.checkip.enable` is `off`, add the SVM name to the `/etc/snapmirror.allow` file.
 - If the `snapmirror.access` option is `legacy` and the `snapmirror.checkip.enable` is `on`, add the IP addresses of the intercluster LIFs to the `/etc/snapmirror.allow` file.
3. Depending on the Data ONTAP version of your 7-Mode system, perform the following steps:
 - a. Allow SnapMirror traffic on all the interfaces by entering the following command:

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

```
options interface.blocked.mgmt_data_traffic off
```
4. If you have set the I2P, read allocations, or NVFAIL options on the volume, perform the following steps:
 - a. Verify that other operations are not impacted if these options are disabled.
 - b. Disable the options by using the following commands:

```
vol options vol_name no_i2p off
```

```
vol options vol_name read_realloc off
```

```
vol options vol_name nvfail off
```

Related information

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

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 setting up intercluster LIFs and verifying network connectivity for transition.

Before you begin

- The cluster and the SVM must already be set up.
For information about setting up the cluster and the SVM, see the *Clustered Data ONTAP Software Setup Guide*.
- The cluster must be reachable by using the cluster-management LIF.
- The cluster must be healthy and none of the nodes must be in takeover mode.
- An SFO policy must be enabled for the aggregates in the SVM to which the volumes will be transitioned.
- The aggregates must be on nodes that have not reached the maximum volume limit.
- If you want to transition volumes from a 32-bit aggregate of a 7-Mode system to a 64-bit aggregate of a cluster, you must have provided an additional 5 percent space in the destination aggregate.

The additional space is required to upgrade the transitioned volume to 64-bit format. For more information about aggregates, see the *Clustered Data ONTAP Physical Storage Management Guide*.

- For establishing an SVM peer relationship when transitioning a volume SnapMirror relationship, the following conditions must be met:
 - The secondary cluster should not have an SVM with the same name as that of the primary SVM.
 - The primary cluster should not have an SVM with the same name as that of the secondary SVM.
 - The name of the source 7-Mode system should not conflict with any of the local SVMs or already peered SVMs.

Step

1. 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
```

10 | 7-Mode Data Transition Using SnapMirror

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

Example

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

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

Example

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

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

Related information

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

Creating a transition peer relationship

You must create a transition peer relationship before you can set up a SnapMirror relationship for transition between a 7-Mode system and a cluster. As a cluster administrator, you can create a transition peer relationship between an SVM and a 7-Mode system by using the `vserver peer transition create` command.

Before you begin

You must have ensured that the name of the source 7-Mode system does not conflict with any of local SVMs or already peered SVMs.

About this task

When creating a transition peer relationship, you can also specify a multipath FQDN or IP address for load balancing the data transfers.

Steps

1. Use the `vserver peer transition create` command to create a transition peer relationship.
2. Use the `vserver peer transition show` to verify that the transition peer relationship is created successfully.

Example of creating and viewing transition peer relationships

The following command creates a transition peer relationship between the SVM vs1.example.com and the 7-Mode system src1.example.com with the multipath address src1-e0d.example.com:

```
cluster1::> vserver peer transition create -local-vserver
vs1.example.com -src-filer-name src1.example.com -multi-path-
address src1-e0d.example.com
```

The following examples show a transition peer relationship between a single SVM (vs1.example.com) and multiple 7-Mode systems:

```
cluster1::> vserver peer transition create -local-vserver
vs1.example.com -src-filer-name
src3.example.com
Transition peering created

cluster1::> vserver peer transition create -local-vserver
vs1.example.com -src-filer-name src2.example.com
Transition peering created

cluster1::> vserver peer transition create -local-vserver
vs1.example.com -src-filer-name src1.example.com
Transition peering created
```

The following output shows the transition peer relationships of the SVM vs1.example.com:

```
cluster1::> vserver peer transition show
Vserver  Source Filer  Multi Path Address
-----  -
vs1.example.com
        src1.example.com
        -
vs1.example.com
        src2.example.com
        -
vs1.example.com
        src3.example.com
```

Transitioning volumes

You can transition a stand-alone volume or volumes that are in data protection relationships (in volume SnapMirror relationships) by using SnapMirror technology.

Choices

- [Transitioning a stand-alone volume](#) on page 12
- [Transitioning a volume SnapMirror relationship](#) on page 17

Related tasks

[Transitioning 7-Mode volumes using SnapMirror](#) on page 4

Transitioning a stand-alone volume

Transitioning a stand-alone volume involves creating a SnapMirror relationship, performing a baseline transfer, performing incremental updates, monitoring the data copy operation, breaking the SnapMirror relationship, and moving client access from the 7-Mode volume to the clustered Data ONTAP volume.

Before you begin

The cluster and SVM must already be set up.

Steps

1. Use the `vol create` command to create a clustered Data ONTAP volume of type `DP` to which the 7-Mode data must be transitioned.

The size of the clustered Data ONTAP volume must be equal to or greater than the size of the 7-Mode volume.

Example

```
cluster1::> volume create -vserver vs1.example.com -volume dst_vol -
aggregate aggr1 -size 20M -type DP
[Job 59] Job succeeded: Successful
```

2. Copy data from the 7-Mode volume to the clustered Data ONTAP volume:
 - a. Use the `snapmirror create` command with the relationship type as `TDP` to create a SnapMirror relationship between the 7-Mode system and the SVM.

Example

```
cluster1::> snapmirror create -source-path system7mode:dataVol20 -
destination-path vs1.example.com:dst_vol -type TDP
Operation succeeded: snapmirror create the relationship with
destination vs1.example.com:dst_vol.
```

- b. Use the `snapmirror initialize` command to start the baseline transfer.

Example

```
cluster1::> snapmirror initialize -destination-path  
vs1.example.com:dst_vol  
Operation is queued: snapmirror initialize of destination  
vs1.example.com:dst_vol.
```

- c. Depending on whether you want to update the clustered Data ONTAP volume manually or by setting up a SnapMirror schedule, perform the appropriate action:

If you want to...	Then...
Update transfers manually	<p data-bbox="436 239 940 269">i. Use the <code>snapmirror update</code> command.</p> <pre data-bbox="483 295 1245 373">cluster1::> snapmirror update -destination-path vs1.example.com:vol_trans</pre> <p data-bbox="436 390 1169 451">ii. Use the <code>snapmirror show</code> command to monitor the data copy status.</p> <pre data-bbox="483 477 1245 1328">cluster1::> snapmirror show -destination-path vs1.example.com:dst_vol Source Path: system7mode:dataVol20 Destination Path: vs1.example.com:dst_vol Relationship Type: TDP SnapMirror Schedule: - Tries Limit: - Throttle (KB/sec): unlimited Mirror State: Snapmirrored Relationship Status: Idle Transfer Snapshot: - Snapshot Progress: - Total Progress: - Snapshot Checkpoint: - Newest Snapshot: vs1.example.com(4053132614)_dst_vol.1 Newest Snapshot Timestamp: 02/13 08:10:46 Exported Snapshot: vs1.example.com(4053132614)_dst_vol.1 Exported Snapshot Timestamp: 02/13 08:10:46 Healthy: true Unhealthy Reason: - Constituent Relationship: false Destination Volume Node: cluster1-01 Relationship ID: e106827a-75b3-11e2- add9-123478563412 Transfer Type: - Transfer Error: - Current Throttle: - Current Transfer Priority: - Last Transfer Type: update Last Transfer Error: - Last Transfer Size: 680KB Last Transfer Duration: 0:0:12 Last Transfer From: system7mode:dataVol20 Last Transfer End Timestamp: 02/13 08:13:40 Progress Last Updated: - Relationship Capability: 8.2 and above Lag Time: 0:8:49 SnapMirror Policy: DPDefault</pre> <p data-bbox="436 1345 739 1374">iii. Go to Step 6 on page 16.</p>

If you want to...	Then...
Perform scheduled update transfers	<p data-bbox="440 244 1147 302">i. Use the <code>job schedule cron create</code> command to create a schedule for update transfers.</p> <pre data-bbox="485 343 1038 392">cluster1::> job schedule cron create -name 15_minute_sched -minute 15</pre> <p data-bbox="440 427 1228 487">ii. Use the <code>snapmirror modify</code> command to apply the schedule to the SnapMirror relationship.</p> <pre data-bbox="485 531 1130 579">cluster1::> snapmirror modify -destination-path vs1.example.com:dst_vol -schedule 15_minute_sched</pre> <p data-bbox="440 614 1166 675">iii. Use the <code>snapmirror show</code> command to monitor the data copy status.</p> <pre data-bbox="485 713 1083 1529">cluster1::> snapmirror show -destination-path vs1.example.com:dst_vol Source Path: system7mode:dataVol120 Destination Path: vs1.example.com:dst_vol Relationship Type: TDP SnapMirror Schedule: 15_minute_sched Tries Limit: - Throttle (KB/sec): unlimited Mirror State: Snapmirrored Relationship Status: Idle Transfer Snapshot: - Snapshot Progress: - Total Progress: - Snapshot Checkpoint: - Newest Snapshot: vs1.example.com(4053132614)_dst_vol.1 Newest Snapshot Timestamp: 02/13 08:10:46 Exported Snapshot: vs1.example.com(4053132614)_dst_vol.1 Exported Snapshot Timestamp: 02/13 08:10:46 Healthy: true Unhealthy Reason: - Constituent Relationship: false Destination Volume Node: cluster1-01 Relationship ID: e106827a-75b3-11e2- add9-123478563412 Transfer Type: - Transfer Error: - Current Throttle: - Current Transfer Priority: - Last Transfer Type: update Last Transfer Error: - Last Transfer Size: 680KB Last Transfer Duration: 0:0:12 Last Transfer From: system7mode:dataVol120 Last Transfer End Timestamp: 02/13 08:13:40 Progress Last Updated: - Relationship Capability: 8.2 and above Lag Time: 0:8:49 SnapMirror Policy: DPDefault</pre>

- 3.** If you have a schedule for incremental transfers, perform the following steps when you are ready to perform cutover:

16 | 7-Mode Data Transition Using SnapMirror

- a. Optional: Use the `snapmirror quiesce` command to disable all future update transfers.

Example

```
cluster1::> snapmirror quiesce -destination-path vs1.example.com:dst_vol
```

- b. Use the `snapmirror modify` command to delete the SnapMirror schedule.

Example

```
cluster1::> snapmirror modify -destination-path vs1.example.com:dst_vol -  
schedule ""
```

- c. Optional: If you quiesced the SnapMirror transfers earlier, use the `snapmirror resume` command to enable SnapMirror transfers.

Example

```
cluster1::> snapmirror resume -destination-path vs1.example.com:dst_vol
```

4. Wait for any ongoing transfers between the 7-Mode volumes and the clustered Data ONTAP volumes to complete, and then disconnect client access from the 7-Mode volumes to start cutover.
5. Use the `snapmirror update` command to perform a final data update to the clustered Data ONTAP volume.

Example

```
cluster1::> snapmirror update -destination-path  
vs1.example.com:dst_vol  
Operation is queued: snapmirror update of destination  
vs1.example.com:dst_vol.
```

6. Use the `snapmirror show` command to verify that the last transfer was successful.
7. Use the `snapmirror break` command to break the SnapMirror relationship between the 7-Mode volume and the clustered Data ONTAP volume.

Example

```
cluster1::> snapmirror break -destination-path vs1.example.com:dst_vol  
[Job 60] Job succeeded: SnapMirror Break Succeeded
```

8. Use the `snapmirror delete` command to delete the SnapMirror relationship between the 7-Mode volume and the clustered Data ONTAP volume.

Example

```
cluster1::> snapmirror delete -destination-path
vs1.example.com:dst_vol
```

9. Use the `snapmirror release` command to remove the SnapMirror relationship information from the 7-Mode system.

Example

```
system7mode> snapmirror release dataVol20 vs1.example.com:dst_vol
```

After you finish

You must delete the SVM peer relationship between the 7-Mode system and the SVM when all the required volumes in the 7-Mode system are transitioned to the SVM.

Related tasks

[Resuming a failed SnapMirror baseline transfer](#) on page 31

Transitioning a volume SnapMirror relationship

You can transition a 7-Mode volume SnapMirror relationship and retain the data protection relationship by transitioning the secondary volume before the primary volume.

Before you begin

- The primary and secondary clusters and SVMs must already be set up.
- For establishing an SVM peer relationship when transitioning a volume SnapMirror relationship, the following conditions must be met:
 - The secondary cluster should not have an SVM with the same name as that of the primary SVM.
 - The primary cluster should not have an SVM with the same name as that of the secondary SVM.

Steps

1. [Transitioning a secondary volume](#) on page 18
2. [Transitioning a primary volume](#) on page 23

Related tasks

[Resuming a failed SnapMirror baseline transfer](#) on page 31

Transitioning a secondary volume

Transitioning a secondary volume involves creating a SnapMirror relationship, performing a baseline transfer, performing incremental updates, and setting up a SnapMirror relationship between the 7-Mode primary volume and the clustered Data ONTAP secondary volume.

Before you begin

The secondary cluster and SVM must already be set up.

Steps

1. Use the `vol create` command to create a clustered Data ONTAP secondary volume of type `DP` to which the 7-Mode data must be transitioned.

The size of the clustered Data ONTAP volume must be equal to or greater than the size of the 7-Mode volume.

Example

```
sec_cluster::> volume create -volume dst_c_vol -aggregate aggr1 -size
20M -type DP
[Job 24] Job succeeded: Successful
```

2. Copy data from the 7-Mode volume to the clustered Data ONTAP volume:
 - a. Use the `snapmirror create` command with the relationship type as `TDP` to create a SnapMirror relationship between the 7-Mode system and the SVM.

Example

```
sec_cluster::> snapmirror create -source-path sec_system:dst_7_vol
-destination-path dst_vserver:dst_c_vol -type TDP
Operation succeeded: snapmirror create the relationship with
destination dst_vserver:dst_c_vol.
```

- b. Use the `snapmirror initialize` command to start the baseline transfer.

Example

```
sec_cluster::> snapmirror initialize -destination-path
dst_vserver:dst_c_vol
Operation is queued: snapmirror initialize of destination
dst_vserver:dst_c_vol.
```

- c. Depending on whether you want to update the clustered Data ONTAP volume manually or by setting up a SnapMirror schedule, perform the appropriate action:

If you want to...	Then...
<p>Update transfers manually</p>	<p>i. Use the <code>snapmirror update</code> command.</p> <pre data-bbox="483 303 1236 373">sec_cluster::> snapmirror update -destination-path dst_vserver:dst_c_vol</pre> <p>ii. Use the <code>snapmirror show</code> command to monitor the data copy status.</p> <pre data-bbox="483 477 1236 1275">sec_cluster::> snapmirror show -destination-path dst_vserver:dst_c_vol Source Path: sec_system:dst_7_vol Destination Path: dst_vserver:dst_c_vol Relationship Type: TDP SnapMirror Schedule: - Tries Limit: - Throttle (KB/sec): unlimited Mirror State: Snapmirrored Relationship Status: Idle Transfer Snapshot: - Snapshot Progress: - Total Progress: - Snapshot Checkpoint: - Newest Snapshot: dst_vserver(4053132614)_dst_c_vol.1 Newest Snapshot Timestamp: 02/13 08:10:46 Exported Snapshot: dst_vserver(4053132614)_dst_c_vol.1 Exported Snapshot Timestamp: 02/13 08:10:46 Healthy: true Unhealthy Reason: - Constituent Relationship: false Destination Volume Node: cluster1-01 Relationship ID: e106827a-75b3-11e2- add9-123478563412 Transfer Type: - Transfer Error: - Current Throttle: - Current Transfer Priority: - Last Transfer Type: update Last Transfer Error: - Last Transfer Size: 680KB Last Transfer Duration: 0:0:12 Last Transfer From: system7mode:dataVol20 Last Transfer End Timestamp: 02/13 08:13:40 Progress Last Updated: - Relationship Capability: 8.2 and above Lag Time: 0:8:49 SnapMirror Policy: DPDefault</pre> <p>iii. Go to Step 6 on page 21.</p>

If you want to...	Then...
Perform scheduled update transfers	<p data-bbox="440 244 1147 302">i. Use the <code>job schedule cron create</code> command to create a schedule for update transfers.</p> <pre data-bbox="485 343 1076 392">sec_cluster::> job schedule cron create -name 15_minute_sched -minute 15</pre> <p data-bbox="440 427 1228 487">ii. Use the <code>snapmirror modify</code> command to apply the schedule to the SnapMirror relationship.</p> <pre data-bbox="485 531 1143 579">sec_cluster::> snapmirror modify -destination-path dst_vserver:dst_c_vol -schedule 15_minute_sched</pre> <p data-bbox="440 614 1166 675">iii. Use the <code>snapmirror show</code> command to monitor the data copy status.</p> <pre data-bbox="485 713 1217 1468">sec_cluster::> snapmirror show -destination-path dst_vserver:dst_c_vol Source Path: sec_system:dst_7_vol Destination Path: dst_vserver:dst_c_vol Relationship Type: TDP SnapMirror Schedule: 15_minute_sched Tries Limit: - Throttle (KB/sec): unlimited Mirror State: Snapmirrored Relationship Status: Idle Transfer Snapshot: - Snapshot Progress: - Total Progress: - Snapshot Checkpoint: - Newest Snapshot: dst_vserver(4053132614)_dst_c_vol.1 Newest Snapshot Timestamp: 02/13 08:10:46 Exported Snapshot: dst_vserver(4053132614)_dst_c_vol.1 Exported Snapshot Timestamp: 02/13 08:10:46 Healthy: true Unhealthy Reason: - Constituent Relationship: false Destination Volume Node: cluster1-01 Relationship ID: e106827a-75b3-11e2- add9-123478563412 Transfer Type: - Transfer Error: - Current Throttle: - Current Transfer Priority: - Last Transfer Type: update Last Transfer Error: - Last Transfer Size: 680KB Last Transfer Duration: 0:0:12 Last Transfer From: system7mode:dataVol120 Last Transfer End Timestamp: 02/13 08:13:40 Progress Last Updated: - Relationship Capability: 8.2 and above Lag Time: 0:8:49 SnapMirror Policy: DPDefault</pre>

3. If you have a schedule for incremental transfers, perform the following steps when you are ready to perform cutover:
 - a. Optional: Use the `snapmirror quiesce` command to disable all future update transfers.

Example

```
sec_cluster::> snapmirror quiesce -destination-path dst_vserver:dst_vol
```

- b. Use the `snapmirror modify` command to delete the SnapMirror schedule.

Example

```
sec_cluster::> snapmirror modify -destination-path dst_vserver:dst_vol -
schedule ""
```

- c. Optional: If you quiesced the SnapMirror transfers earlier, use the `snapmirror resume` command to enable SnapMirror transfers.

Example

```
sec_cluster::> snapmirror resume -destination-path dst_vserver:dst_vol
```

4. Wait for any ongoing transfers between the 7-Mode volumes and the clustered Data ONTAP volumes to complete, and then disconnect client access from the 7-Mode volumes to start cutover.
5. Use the `snapmirror update` command to perform a final data update to the clustered Data ONTAP volume.

Example

```
sec_cluster::> snapmirror update -destination-path dst_vserver:dst_vol
Operation is queued: snapmirror update of destination
dst_vserver:dst_vol.
```

6. Use the `snapmirror show` command to verify that the last transfer was successful.
7. Use the `snapmirror break` command to break the SnapMirror relationship between the 7-Mode secondary volume and the clustered Data ONTAP secondary volume.

Example

```
sec_cluster::> snapmirror break -destination-path dst_vserver:dst_vol
[Job 60] Job succeeded: SnapMirror Break Succeeded
```

8. Use the `snapmirror delete` command to delete the SnapMirror relationship between the 7-Mode secondary volume and the clustered Data ONTAP secondary volume.

Example

```
sec_cluster::> snapmirror delete -destination-path dst_vserver:dst_vol
```

22 | 7-Mode Data Transition Using SnapMirror

9. Use the `snapmirror release` command to remove the SnapMirror relationship information from the 7-Mode system.

Example

```
system7mode> snapmirror release dataVol20 vs1.example.com:dst_vol
```

10. Establish a disaster recovery relationship between the 7-Mode primary volume and clustered Data ONTAP secondary volume:
 - a. Use the `vserver peer transition create` command to create an SVM peer relationship between the 7-Mode primary volume and the clustered Data ONTAP secondary volume.

Example

```
sec_cluster::> vserver peer transition create -local-vserver  
dst_vserver -src-filer-name src_system  
Transition peering created
```

- b. Use the `job schedule cron create` command to create a job schedule that matches the schedule configured for the 7-Mode SnapMirror relationship.

Example

```
sec_cluster::> job schedule cron create -name 15_minute_sched -  
minute 15
```

- c. Use the `snapmirror create` command to create a SnapMirror relationship between the 7-Mode primary volume and the clustered Data ONTAP secondary volume.

Example

```
sec_cluster::> snapmirror create -source-path src_system:src_7_vol  
-destination-path dst_vserver:dst_c_vol -type TDP -schedule  
15_minute_sched  
Operation succeeded: snapmirror create the relationship with  
destination dst_vserver:dst_c_vol.
```

- d. Use the `snapmirror resync` command to resynchronize the clustered Data ONTAP secondary volume.

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

Example

```
sec_cluster::> snapmirror resync -destination-path
dst_vserver:dst_c_vol
```

After you finish

- Delete the SVM peer relationship between the secondary 7-Mode system and the secondary SVM when all the required volumes in the 7-Mode system are transitioned to the SVM.
- Delete the SnapMirror relationship between the 7-Mode primary and the 7-Mode secondary systems.

Transitioning a primary volume

Transitioning a primary volume involves copying data from the 7-Mode primary volumes to the clustered Data ONTAP primary volumes, deleting the disaster recovery relationship between the 7-Mode primary and clustered Data ONTAP secondary volumes, and establishing a SnapMirror relationship between the clustered Data ONTAP primary and secondary volumes.

Before you begin

The primary cluster and SVM must already be set up.

Steps

1. Use the `vol create` command to create a clustered Data ONTAP volume of type `DP` to which the 7-Mode data must be transitioned.

The size of the clustered Data ONTAP volume must be equal to or greater than the size of the 7-Mode volume.

Example

```
pri_cluster::> vol create -vserver src_vserver -volume src_c_vol -
aggregate aggr1 -size 50M -state online -type DP -policy default
(volume create)
[Job 1480] Job succeeded: Successful
```

2. Copy the data from the 7-Mode primary volume to the clustered Data ONTAP primary volume:
 - a. Use the `snapmirror create` command with the relationship type as `TDP` to create a SnapMirror relationship between the 7-Mode system and the SVM.

Example

```
pri_cluster::> snapmirror create -source-path src_system:finance -
destination-path src_vserver:src_c_vol -type TDP
```

24 | 7-Mode Data Transition Using SnapMirror

```
Operation succeeded: snapmirror create the relationship with
destination src_vserver:src_c_vol.
```

- b. Use the `snapmirror initialize` command to start the baseline transfer.

Example

```
pri_cluster::> snapmirror initialize -destination-path
src_vserver:src_c_vol
Operation is queued: snapmirror initialize of destination
src_vserver:src_c_vol.
```

- c. Depending on whether you want to update the clustered Data ONTAP volume manually or by setting up a SnapMirror schedule, perform the appropriate action:

If you want to...	Then...
Update transfers manually	<p data-bbox="440 244 942 270">i. Use the <code>snapmirror update</code> command.</p> <pre data-bbox="494 309 1143 361">pri_cluster::> snapmirror update -destination-path src_vserver:src_c_vol</pre> <p data-bbox="440 395 1166 453">ii. Use the <code>snapmirror show</code> command to monitor the data copy status.</p> <pre data-bbox="494 496 1220 1269">pri_cluster::> snapmirror show -destination-path src_vserver:src_c_vol Source Path: pri_system:src_7_vol Destination Path: src_vserver:src_c_vol Relationship Type: TDP SnapMirror Schedule: - Tries Limit: - Throttle (KB/sec): unlimited Mirror State: Snapmirrored Relationship Status: Idle Transfer Snapshot: - Snapshot Progress: - Total Progress: - Snapshot Checkpoint: - Newest Snapshot: src_vserver(4053132614)_src_c_vol.1 Newest Snapshot Timestamp: 02/13 08:10:46 Exported Snapshot: src_vserver(4053132614)_src_c_vol.1 Exported Snapshot Timestamp: 02/13 08:10:46 Healthy: true Unhealthy Reason: - Constituent Relationship: false Destination Volume Node: cluster1-01 Relationship ID: e106827a-75b3-11e2- add9-123478563412 Transfer Type: - Transfer Error: - Current Throttle: - Current Transfer Priority: - Last Transfer Type: update Last Transfer Error: - Last Transfer Size: 680KB Last Transfer Duration: 0:0:12 Last Transfer From: system7mode:dataVol20 Last Transfer End Timestamp: 02/13 08:13:40 Progress Last Updated: - Relationship Capability: 8.2 and above Lag Time: 0:8:49 SnapMirror Policy: DPDefault</pre> <p data-bbox="440 1307 626 1333">iii. Go to Step 6.</p>

If you want to...	Then...
Perform scheduled update transfers	<p data-bbox="440 244 1147 302">i. Use the <code>job schedule cron create</code> command to create a schedule for update transfers.</p> <pre data-bbox="485 343 1076 392">pri_cluster::> job schedule cron create -name 15_minute_sched -minute 15</pre> <p data-bbox="440 427 1228 487">ii. Use the <code>snapmirror modify</code> command to apply the schedule to the SnapMirror relationship.</p> <pre data-bbox="485 531 1143 579">pri_cluster::> snapmirror modify -destination-path src_vserver:src_c_vol -schedule 15_minute_sched</pre> <p data-bbox="440 614 1166 675">iii. Use the <code>snapmirror show</code> command to monitor the data copy status.</p> <pre data-bbox="485 713 1217 1468">pri_cluster::> snapmirror show -destination-path src_vserver:src_c_vol Source Path: pri_system:src_7_vol Destination Path: src_vserver:src_c_vol Relationship Type: TDP SnapMirror Schedule: 15_minute_sched Tries Limit: - Throttle (KB/sec): unlimited Mirror State: Snapmirrored Relationship Status: Idle Transfer Snapshot: - Snapshot Progress: - Total Progress: - Snapshot Checkpoint: - Newest Snapshot: src_vserver(4053132614)_src_c_vol.1 Newest Snapshot Timestamp: 02/13 08:10:46 Exported Snapshot: src_vserver(4053132614)_src_c_vol.1 Exported Snapshot Timestamp: 02/13 08:10:46 Healthy: true Unhealthy Reason: - Constituent Relationship: false Destination Volume Node: cluster1-01 Relationship ID: e106827a-75b3-11e2- add9-123478563412 Transfer Type: - Transfer Error: - Current Throttle: - Current Transfer Priority: - Last Transfer Type: update Last Transfer Error: - Last Transfer Size: 680KB Last Transfer Duration: 0:0:12 Last Transfer From: system7mode:dataVol120 Last Transfer End Timestamp: 02/13 08:13:40 Progress Last Updated: - Relationship Capability: 8.2 and above Lag Time: 0:8:49 SnapMirror Policy: DPDefault</pre>

3. If you have a schedule for incremental transfers, perform the following steps when you are ready to perform cutover:
 - a. Optional: Use the `snapmirror quiesce` command to disable all future update transfers.

Example

```
pri_cluster::> snapmirror quiesce -destination-path src_vserver:src_c_vol
```

- b. Use the `snapmirror modify` command to delete the SnapMirror schedule.

Example

```
pri_cluster::> snapmirror modify -destination-path src_vserver:src_c_vol -
schedule ""
```

- c. Optional: If you quiesced the SnapMirror transfers earlier, use the `snapmirror resume` command to enable SnapMirror transfers.

Example

```
pri_cluster::> snapmirror resume -destination-path src_vserver:src_c_vol
```

4. Create an SVM peer relationship between the clustered Data ONTAP secondary and primary SVMs.

- a. Use the `cluster peer create` command to create a cluster peer relationship.

Example

```
pri_cluster::> cluster peer create -peer-addr sec_cluster -
username admin
Password: *****
```

- b. From the source cluster, use the `vserver peer create` command to create a SVM peer relationship between the clustered Data ONTAP primary and secondary volumes.

Example

```
pri_cluster::> vserver peer create -vserver src_vserver -
peervserver src_c_vserver -applications snapmirror -peer-cluster
sec_cluster
```

- c. From the destination cluster, use the `vserver peer accept` command to accept the SVM peer request and establish the SVM peer relationship.

Example

```
sec_cluster::> vserver peer accept -vserver dst_vserver -
peervserver src_vserver
```

5. From the destination cluster, use the `snapmirror quiesce` command to suspend any data transfers between the 7-Mode primary volume and the clustered Data ONTAP secondary volume if a schedule is set up for update transfers.

Example

```
sec_cluster::> snapmirror quiesce -destination-path dst_vserver:dst_c_vol
```

6. Monitor the data copy operation and initiate cutover:
 - a. Wait for any ongoing transfers from the 7-Mode primary volumes to the clustered Data ONTAP primary and clustered Data ONTAP secondary volumes to complete, and disconnect client access from the 7-Mode primary volume to start cutover.
 - b. Use the `snapmirror update` command to perform a final data update to the clustered Data ONTAP primary volume from the 7-Mode primary volume.

Example

```
pri_cluster::> snapmirror update -destination-path  
src_vserver:src_c_vol
```

- c. Use the `snapmirror break` command to break the SnapMirror relationship between the 7-Mode primary volume and clustered Data ONTAP primary volume.

Example

```
pri_cluster::> snapmirror break -destination-path  
src_vserver:src_c_vol  
[Job 1485] Job is queued: snapmirror break for destination  
src_vserver:src_c_vol.
```

- d. Use the `snapmirror delete` command to delete the relationship.

Example

```
pri_cluster::> snapmirror delete -destination-path  
src_vserver:src_c_vol
```

- e. Use the `snapmirror release` command to remove the SnapMirror relationship information from the 7-Mode system.

Example

```
system7mode> snapmirror release dataVol20 vs1.example.com:dst_vol
```

7. From the destination cluster, break and delete the disaster recovery relationship between the 7-Mode primary volume and clustered Data ONTAP secondary volume.

- a. Use the `snapmirror break` command to break the disaster recovery relationship between the 7-Mode primary volume and clustered Data ONTAP secondary volume.

Example

```
sec_cluster::> snapmirror break -destination-path
dst_vserver:dst_c_vol
[Job 1485] Job is queued: snapmirror break for destination
dst_vserver:dst_c_vol.
```

- b. Use the `snapmirror delete` command to delete the relationship.

Example

```
sec_cluster::> snapmirror delete -destination-path
dst_vserver:dst_c_vol
```

- c. Use the `snapmirror release` command to remove the SnapMirror relationship information from the 7-Mode system.

Example

```
system7mode> snapmirror release dataVol20 vs1.example.com:dst_vol
```

8. From the destination cluster, establish a SnapMirror relationship between the clustered Data ONTAP primary and secondary volumes:
 - a. Use the `snapmirror create` command to create a SnapMirror relationship between the clustered Data ONTAP primary and secondary volumes.

Example

```
sec_cluster::> snapmirror create -source-path
src_vserver:src_c_vol -destination-path dst_vserver:dst_c_vol -
type DP -schedule 15_minute_sched
```

- b. Use the `snapmirror resync` command to resynchronize the SnapMirror relationship between the clustered Data ONTAP volumes.

For successful resynchronization, a common Snapshot copy must exist between the clustered Data ONTAP primary and secondary volumes.

Example

```
sec_cluster::> snapmirror resync -destination-path
dst_vserver:dst_c_vol
```

30 | 7-Mode Data Transition Using SnapMirror

- c. Use the `snapmirror show` command to verify that the status of SnapMirror resynchronization shows `SnapMirrored`.

Note: You must ensure that the SnapMirror resynchronization is successful to make the clustered Data ONTAP secondary volume available for read-only access.

After you finish

You must delete the SVM peer relationship between the 7-Mode system and the SVM when all the required volumes in the 7-Mode system are transitioned to the SVM.

Resuming a failed SnapMirror baseline transfer

During transition, SnapMirror baseline transfers can fail due to a number of reasons, such as loss of network connectivity, transfer aborted, or controller failover. After rectifying the cause of failure, you can resume the SnapMirror transfers if a restart checkpoint is available.

About this task

If the restart checkpoint for the baseline transfer is not available, you must delete and re-create the volume, reestablish the SnapMirror relationship, and initiate the transition again.

Steps

1. From the destination cluster, use the `snapmirror show` command with the `-snapshot-checkpoint` parameter to view the status of the baseline transfer and the restart checkpoint.

Example

```
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. If the SnapMirror checkpoint exists, use the `snapmirror initialize` command to resume the baseline transfer.

Example

```
cluster2::> snapmirror initialize -destination-path dest_vserver:vol3
```

Recovering from a disaster at the 7-Mode site during transition

If you have established a SnapMirror disaster recovery (DR) relationship between the 7-Mode primary volume and the clustered Data ONTAP secondary volume and if a disaster occurs at the 7-Mode primary site, you can direct client access to the clustered Data ONTAP secondary volume. After the 7-Mode primary volume comes back online, you have to perform additional steps to redirect the clients to the clustered Data ONTAP primary volume.

About this task

To retain any data written on the clustered Data ONTAP secondary volume after the disaster, you must transition the 7-Mode primary volume after the 7-Mode primary volume comes back online and establish a SnapMirror relationship between the clustered Data ONTAP primary and secondary volumes. You can then redirect the clients to the clustered Data ONTAP primary volume.

SnapMirror resynchronization from clustered Data ONTAP volumes to the 7-Mode volumes is not supported. Therefore, if you reestablish the DR relationship between the 7-Mode primary volume and clustered Data ONTAP secondary volume after the disaster, any data written on the secondary clustered Data ONTAP will be lost.

Steps

1. [Redirecting clients to the clustered Data ONTAP secondary volume after disaster](#) on page 33
When a disaster strikes at the 7-Mode primary volume, you can redirect the clients from the 7-Mode primary volume to the clustered Data ONTAP secondary volume.
2. [Transitioning the 7-Mode primary as a stand-alone volume](#) on page 33
After the 7-Mode primary volume comes back online after the disaster, you must transition the 7-Mode primary volume. Because all SnapMirror relationships to the 7-Mode primary volume are broken and deleted at this stage, you can transition a stand-alone volume for this type of transition.
3. [Redirecting clients to the clustered Data ONTAP primary volume](#) on page 38
After transition to the clustered Data ONTAP primary volume is complete, you can resynchronize the clustered Data ONTAP primary volume for the data written on the clustered Data ONTAP secondary volume. You can then redirect the clients to the clustered Data ONTAP primary volume.

Redirecting clients to the clustered Data ONTAP secondary volume after a disaster

If you have established a SnapMirror disaster recovery (DR) relationship between the 7-Mode primary volume and the clustered Data ONTAP secondary volume and if a disaster occurs at the 7-Mode primary site, you must redirect client access to the clustered Data ONTAP secondary volume.

Steps

1. From the secondary cluster, use the `snapmirror break` command to break the SnapMirror relationship between the 7-Mode primary volume and the clustered Data ONTAP secondary volume.

Example

```
sec_cluster::> snapmirror break -destination-path  
dst_vserver:dst_c_vol
```

2. From the secondary cluster, use the `snapmirror delete` command to delete the SnapMirror relationship between the 7-Mode primary volume and the clustered Data ONTAP secondary volume.

Example

```
sec_cluster::> snapmirror delete -destination-path  
dst_vserver:dst_c_vol
```

3. Redirect client access to the clustered Data ONTAP secondary volume.

For more information about setting up client access in clustered Data ONTAP, see the *Clustered Data ONTAP File Access and Protocols Management Guide*.

Transitioning a stand-alone volume

After the 7-Mode primary volume comes back online after the disaster, you must transition the 7-Mode primary volume. Because all SnapMirror relationships to the 7-Mode primary volume are broken and deleted at this stage, you can transition a stand-alone volume for this type of transition.

Before you begin

The cluster and SVM must already be set up.

Steps

1. Use the `vol create` command to create a clustered Data ONTAP volume of type DP to which the 7-Mode data must be transitioned.

The size of the clustered Data ONTAP volume must be equal to or greater than the size of the 7-Mode volume.

Example

```
cluster1::> volume create -vserver vs1.example.com -volume dst_vol -  
aggregate aggr1 -size 20M -type DP  
[Job 59] Job succeeded: Successful
```

2. Copy data from the 7-Mode volume to the clustered Data ONTAP volume:
 - a. Use the `snapmirror create` command with the relationship type as TDP to create a SnapMirror relationship between the 7-Mode system and the SVM.

Example

```
cluster1::> snapmirror create -source-path system7mode:dataVol20 -  
destination-path vs1.example.com:dst_vol -type TDP  
Operation succeeded: snapmirror create the relationship with  
destination vs1.example.com:dst_vol.
```

- b. Use the `snapmirror initialize` command to start the baseline transfer.

Example

```
cluster1::> snapmirror initialize -destination-path  
vs1.example.com:dst_vol  
Operation is queued: snapmirror initialize of destination  
vs1.example.com:dst_vol.
```

- c. Depending on whether you want to update the clustered Data ONTAP volume manually or by setting up a SnapMirror schedule, perform the appropriate action:

If you want to...	Then...
Update transfers manually	<p data-bbox="440 244 942 270">i. Use the <code>snapmirror update</code> command.</p> <pre data-bbox="485 296 1103 357">cluster1::> snapmirror update -destination-path vs1.example.com:vol_trans</pre> <p data-bbox="440 395 1166 456">ii. Use the <code>snapmirror show</code> command to monitor the data copy status.</p> <pre data-bbox="485 496 1083 1312">cluster1::> snapmirror show -destination-path vs1.example.com:dst_vol Source Path: system7mode:dataVol20 Destination Path: vs1.example.com:dst_vol Relationship Type: TDP SnapMirror Schedule: - Tries Limit: - Throttle (KB/sec): unlimited Mirror State: Snapmirrored Relationship Status: Idle Transfer Snapshot: - Snapshot Progress: - Total Progress: - Snapshot Checkpoint: - Newest Snapshot: vs1.example.com(4053132614)_dst_vol.1 Newest Snapshot Timestamp: 02/13 08:10:46 Exported Snapshot: vs1.example.com(4053132614)_dst_vol.1 Exported Snapshot Timestamp: 02/13 08:10:46 Healthy: true Unhealthy Reason: - Constituent Relationship: false Destination Volume Node: cluster1-01 Relationship ID: e106827a-75b3-11e2- add9-123478563412 Transfer Type: - Transfer Error: - Current Throttle: - Current Transfer Priority: - Last Transfer Type: update Last Transfer Error: - Last Transfer Size: 680KB Last Transfer Duration: 0:0:12 Last Transfer From: system7mode:dataVol20 Last Transfer End Timestamp: 02/13 08:13:40 Progress Last Updated: - Relationship Capability: 8.2 and above Lag Time: 0:8:49 SnapMirror Policy: DPDefault</pre> <p data-bbox="440 1347 747 1373">iii. Go to Step 6 on page 16.</p>

If you want to...	Then...
Perform scheduled update transfers	<p data-bbox="440 244 1147 302">i. Use the <code>job schedule cron create</code> command to create a schedule for update transfers.</p> <pre data-bbox="485 343 1038 392">cluster1::> job schedule cron create -name 15_minute_sched -minute 15</pre> <p data-bbox="440 427 1228 487">ii. Use the <code>snapmirror modify</code> command to apply the schedule to the SnapMirror relationship.</p> <pre data-bbox="485 531 1130 579">cluster1::> snapmirror modify -destination-path vs1.example.com:dst_vol -schedule 15_minute_sched</pre> <p data-bbox="440 612 1166 673">iii. Use the <code>snapmirror show</code> command to monitor the data copy status.</p> <pre data-bbox="485 713 1083 1529">cluster1::> snapmirror show -destination-path vs1.example.com:dst_vol Source Path: system7mode:dataVol120 Destination Path: vs1.example.com:dst_vol Relationship Type: TDP SnapMirror Schedule: 15_minute_sched Tries Limit: - Throttle (KB/sec): unlimited Mirror State: Snapmirrored Relationship Status: Idle Transfer Snapshot: - Snapshot Progress: - Total Progress: - Snapshot Checkpoint: - Newest Snapshot: vs1.example.com(4053132614)_dst_vol.1 Newest Snapshot Timestamp: 02/13 08:10:46 Exported Snapshot: vs1.example.com(4053132614)_dst_vol.1 Exported Snapshot Timestamp: 02/13 08:10:46 Healthy: true Unhealthy Reason: - Constituent Relationship: false Destination Volume Node: cluster1-01 Relationship ID: e106827a-75b3-11e2- add9-123478563412 Transfer Type: - Transfer Error: - Current Throttle: - Current Transfer Priority: - Last Transfer Type: update Last Transfer Error: - Last Transfer Size: 680KB Last Transfer Duration: 0:0:12 Last Transfer From: system7mode:dataVol120 Last Transfer End Timestamp: 02/13 08:13:40 Progress Last Updated: - Relationship Capability: 8.2 and above Lag Time: 0:8:49 SnapMirror Policy: DPDefault</pre>

- 3.** If you have a schedule for incremental transfers, perform the following steps when you are ready to perform cutover:

- a. Optional: Use the `snapmirror quiesce` command to disable all future update transfers.

Example

```
cluster1::> snapmirror quiesce -destination-path vs1.example.com:dst_vol
```

- b. Use the `snapmirror modify` command to delete the SnapMirror schedule.

Example

```
cluster1::> snapmirror modify -destination-path vs1.example.com:dst_vol -
schedule ""
```

- c. Optional: If you quiesced the SnapMirror transfers earlier, use the `snapmirror resume` command to enable SnapMirror transfers.

Example

```
cluster1::> snapmirror resume -destination-path vs1.example.com:dst_vol
```

4. Wait for any ongoing transfers between the 7-Mode volumes and the clustered Data ONTAP volumes to complete, and then disconnect client access from the 7-Mode volumes to start cutover.
5. Use the `snapmirror update` command to perform a final data update to the clustered Data ONTAP volume.

Example

```
cluster1::> snapmirror update -destination-path
vs1.example.com:dst_vol
Operation is queued: snapmirror update of destination
vs1.example.com:dst_vol.
```

6. Use the `snapmirror show` command to verify that the last transfer was successful.
7. Use the `snapmirror break` command to break the SnapMirror relationship between the 7-Mode volume and the clustered Data ONTAP volume.

Example

```
cluster1::> snapmirror break -destination-path vs1.example.com:dst_vol
[Job 60] Job succeeded: SnapMirror Break Succeeded
```

8. Use the `snapmirror delete` command to delete the SnapMirror relationship between the 7-Mode volume and the clustered Data ONTAP volume.

Example

```
cluster1::> snapmirror delete -destination-path
vs1.example.com:dst_vol
```

9. Use the `snapmirror release` command to remove the SnapMirror relationship information from the 7-Mode system.

Example

```
system7mode> snapmirror release dataVol20 vs1.example.com:dst_vol
```

After you finish

You must delete the SVM peer relationship between the 7-Mode system and the SVM when all the required volumes in the 7-Mode system are transitioned to the SVM.

Redirecting clients to the clustered Data ONTAP primary volume

After the 7-Mode primary volume comes back online, you can transition the 7-Mode primary volume, establish a SnapMirror relationship with the clustered Data ONTAP secondary volume, and redirect client access to the clustered Data ONTAP primary volume.

Steps

1. Create the SVM peer relationship between the primary and secondary SVMs.
 - a. Use the `cluster peer create` command to create the cluster peer relationship.

Example

```
pri_cluster::> cluster peer create -peer-addr sec_cluster -
username admin
Password: *****
```

- b. From the source cluster, use the `vserver peer create` command to create an SVM peer relationship between the clustered Data ONTAP primary volume and clustered Data ONTAP secondary volume.

Example

```
pri_cluster::> vserver peer create -vserver src_vserver -
peer_vserver src_c_vserver -applications snapmirror -peer-cluster
sec_cluster
```

- c. From the destination cluster, use the `vserver peer accept` command to accept the SVM peer request and establish the SVM peer relationship.

Example

```
sec_cluster::> vserver peer accept -vserver dst_vserver -
peervserver src_vserver
```

2. Use the `snapmirror create` command to create a SnapMirror relationship with the clustered Data ONTAP secondary volume as the source and the clustered Data ONTAP primary volume as destination.

Example

```
pri_cluster::> snapmirror create -source-path dst_vserver:dst_c_vol -
destination-path src_vserver:src_c_vol
```

3. From the primary cluster, use the `snapmirror resync` command to resynchronize the clustered Data ONTAP secondary volume.

Example

```
pri_cluster::> snapmirror resync -source-path dst_vserver:dst_c_vol -
destination-path src_vserver:src_c_vol
```

Wait till the resynchronization finishes. The SnapMirror state changes to SnapMirrored when resynchronization is complete.

4. When you are ready to switch over to the clustered Data ONTAP primary volume, disconnect client access from the clustered Data ONTAP secondary volume.
5. From the primary cluster, use the `snapmirror update` command to update the primary volume.

Example

```
pri_cluster::> snapmirror update -destination-path
src_vserver:src_c_vol
```

6. From the primary cluster, use the `snapmirror break` command to break the SnapMirror relationship between the clustered Data ONTAP primary and secondary volumes.

Example

```
pri_cluster::> snapmirror break -destination-path
src_vserver:src_c_vol
```

7. Enable client access to the clustered Data ONTAP primary volume.
8. From the primary cluster, use the `snapmirror delete` command to delete the SnapMirror relationship between the clustered Data ONTAP primary and secondary volumes.

Example

```
pri_cluster::> snapmirror delete -destination-path  
src_vserver:src_c_vol
```

9. From the secondary cluster, use the `snapmirror create` command to create a SnapMirror relationship with the clustered Data ONTAP primary volume as the source and the clustered Data ONTAP secondary volume as destination, with a schedule similar to the previous schedule between the 7-Mode primary volume and clustered Data ONTAP secondary volume.

Example

```
sec_cluster::> snapmirror create -source-path src_vserver:src_c_vol -  
destination-path dst_vserver:dst_c_vol -schedule 15_minute_sched
```

10. From the secondary cluster, use the `snapmirror resync` command to resynchronize the clustered Data ONTAP primary volume.

Example

```
sec_cluster::> snapmirror resync -source-path src_vserver:src_c_vol -  
destination-path dst_vserver:dst_c_vol
```

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 transition using SnapMirror [4](#)
- preparing for transition [7](#)

7-Mode versions

- supported for transition to clustered Data ONTAP [6](#)

B

baseline transfers

- resuming [31](#)

C

clusters

- preparing for transition [9](#)

comments

- how to send feedback about documentation [43](#)

configuration

- not supported for transition [5](#)

creating

- transition peer relationships [10](#)

D

Data ONTAP

- 7-Mode versions supported for transition [6](#)

disaster recovery

- at the 7-Mode site [32](#)
- redirecting clients to the clustered Data ONTAP primary volume [38](#)
- redirecting clients to the clustered Data ONTAP secondary volume [33](#)

doccomments

- how to send feedback about documentation by using [43](#)

documentation

- how to send feedback about [43](#)

F

features

- not supported for transition [5](#)

feedback

- how to send comments about documentation [43](#)

H

high-quality information

- how to send feedback about improving documentation [43](#)

I

information

- how to send feedback about improving documentation [43](#)

L

licenses

- requirements for transition [7](#)

M

multiple paths

- considerations for data copy using [6](#)

P

peers

- creating relationship for transition [10](#)

planning

- to transition 7-Mode volumes [4](#)

preparing

- 7-Mode system for transition [7](#)
- cluster for transition [9](#)
- for transition [7](#)

primary volumes

- transitioning [23](#)

Q

quality documentation

- how to send feedback about improving [43](#)

R

requirements

- 7-Mode versions supported for transition [6](#)
- licenses for transition [7](#)

S

SAN

- unsupported features [5](#)

secondary volumes

- transitioning [18](#)

SnapMirror

- transitioning volumes [11](#)

- when to use for transition [4](#)

- workflow for transitioning volumes [4](#)

SnapMirror relationships

- considerations for transition [6](#)

- transitioning [17](#)

stand-alone volumes

- transitioning [12](#), [33](#)

suggestions

- how to send feedback about documentation [43](#)

T

transfers

- transition considerations for concurrent SnapMirror [6](#)

transition

- 7-Mode versions supported [6](#)

- considerations for SnapMirror [6](#)

- license requirements [7](#)

- planning for [4](#)

- preparing the 7-Mode system [7](#)

- preparing the cluster for [9](#)

- preparing to [7](#)

- recovering from a disaster [32](#)

- redirecting clients to the clustered Data ONTAP

- primary volume [38](#)

- redirecting clients to the clustered Data ONTAP

- secondary volume [33](#)

- resuming a failed baseline transfer [31](#)

- unsupported features [5](#)

- using SnapMirror [4](#)

- volume SnapMirror relationships [17](#)

transition peer relationships

- creating [10](#)

transitioning

- primary volumes [23](#)

- secondary volumes [18](#)

- stand-alone volumes [12](#), [33](#)

- volumes using SnapMirror [11](#)

U

unsupported features

- for transition [5](#)

V

volume SnapMirror relationships

- transitioning [17](#)

volume transition

- planning for [4](#)

volumes

- transitioning primary volumes [23](#)

- transitioning secondary volumes [18](#)

- transitioning stand-alone volumes [12](#), [33](#)

- transitioning using SnapMirror [11](#)

W

workflow

- for transitioning 7-Mode volumes using SnapMirror [4](#)