



Completing the transition

ONTAP 7-Mode Transition

NetApp

February 11, 2024

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

Table of Contents

- Completing the transition 1
 - Restrictions during preproduction testing 1
 - Rehosting transitioned volumes to a different SVM 3
 - Verifying the transitioned configurations 6
 - Performing manual configuration tasks after transition 6
 - Testing the workloads and applications 7
 - Committing the copy-free transition project 8


Completing the transition


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

Restrictions during preproduction testing

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

Operations that are blocked

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

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

Operations that are not recommended

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

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

Related information

[ONTAP 9 commands](#)

Rehosting transitioned volumes to a different SVM

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

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

Rehosting is a disruptive operation.

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

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

Steps

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

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

• Rehosting NFS volumes

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

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

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

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

NFS management

• Rehosting CIFS volumes

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

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

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

```
volume mount
```

- f. Create CIFS shares for the rehosted volume:

```
vserver cifs share create
```

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

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

[SMB/CIFS management](#)

• Rehosting volumes in SnapMirror relationships

- a. Record the SnapMirror relationship type:

```
snapmirror show
```

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

```
snapmirror delete
```



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

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

```
snapmirror release
```

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

- d. Switch to the advanced privilege level:

```
set -privilege advanced
```

- e. Rehost the volume on the destination SVM:

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

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

```
vserver peer create
```

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

```
snapmirror create
```

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

- h. Resynchronize the data protection relationship:

```
snapmirror resync
```

[Data protection](#)

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

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

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

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

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

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



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

Related information

[ONTAP 9 commands](#)

Verifying the transitioned configurations

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

Steps

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

Performing manual configuration tasks after transition

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

Steps

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

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

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

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

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

[Transitioning a SnapMirror relationship](#)

Testing the workloads and applications

You should manually test all workloads and applications in the preproduction environment. You can then start production for a brief duration to evaluate if rollback is required before committing the project.

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



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

Some operations are restricted during preproduction testing.

Restrictions during preproduction testing

Steps

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

SAN host transition and remediation

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

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

- a. Log in to the advanced privilege level:

```
set -privilege advanced
```

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

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

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

You can start serving production data.



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

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

Related information

Performing a transition rollback to 7-Mode

Committing the copy-free transition project

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

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

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

Steps

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

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

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.