



Copy-free transition workflow

ONTAP 7-Mode Transition

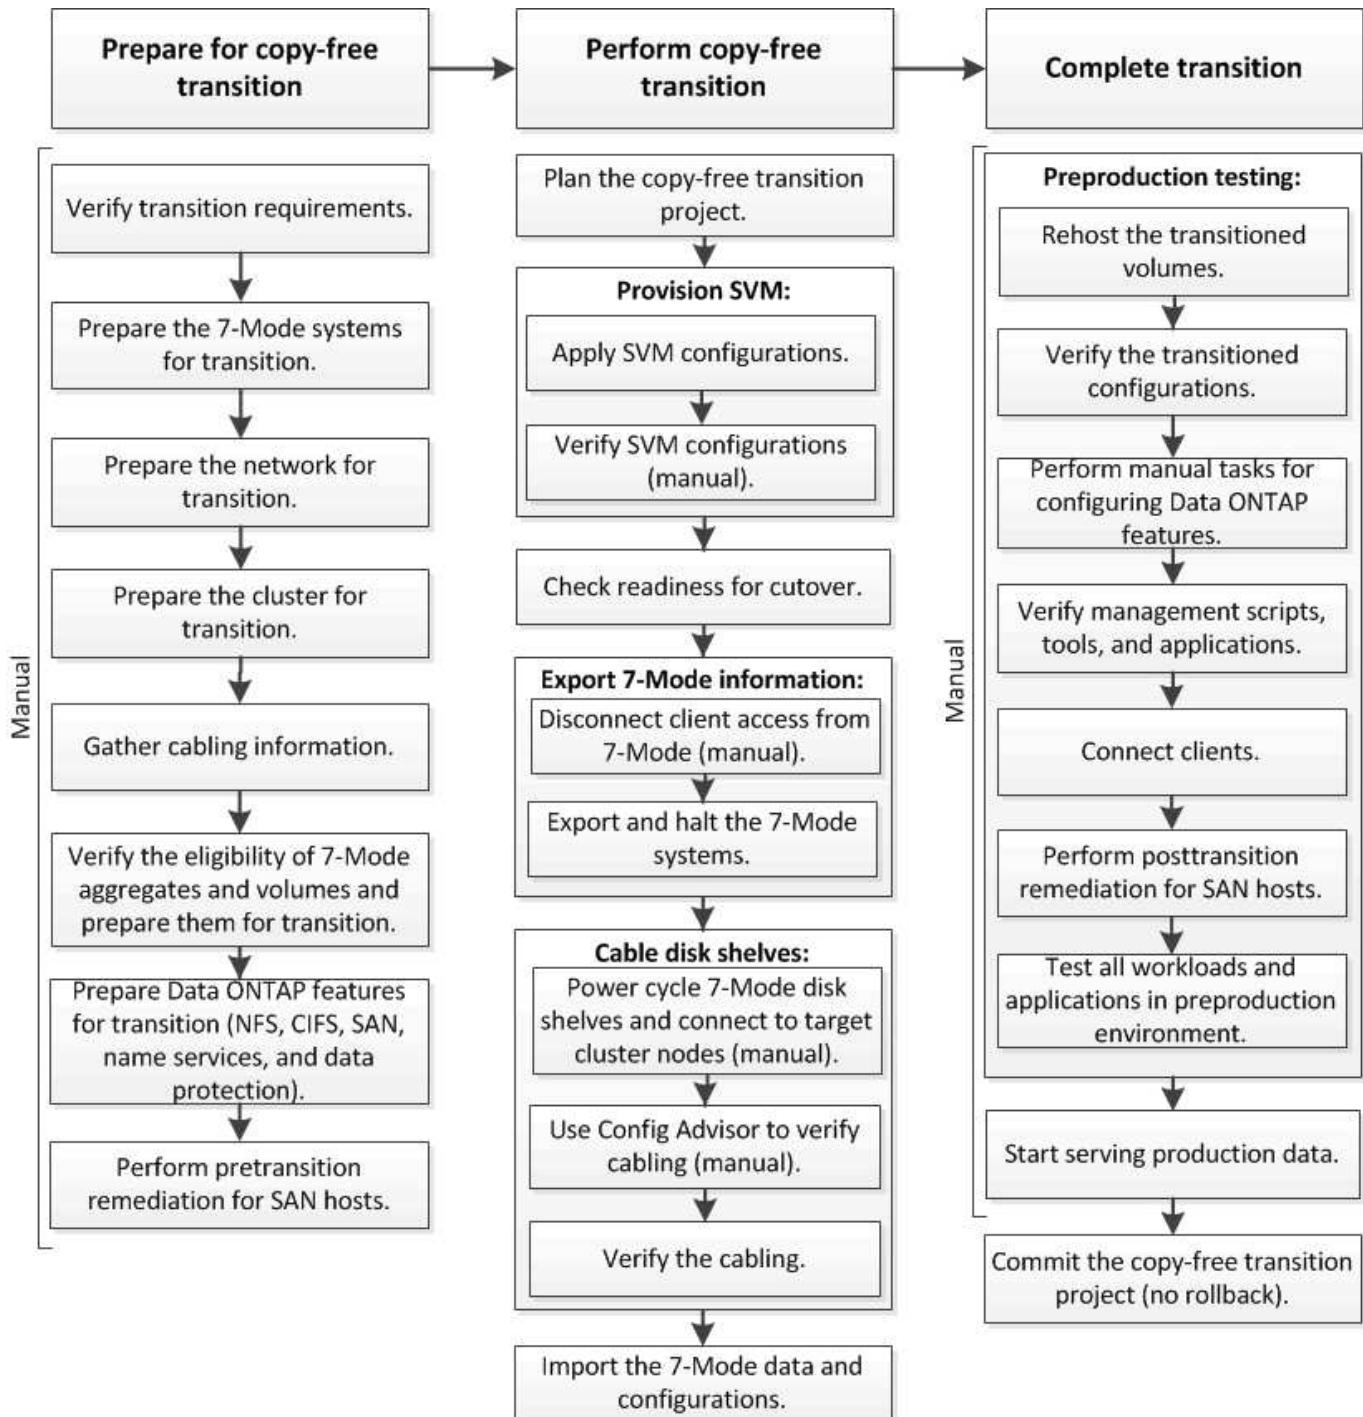
NetApp
October 09, 2025

Table of Contents

Copy-free transition workflow	1
Phases of copy-free transition	2
Planning the project	2
SVM provisioning	3
Exporting storage configurations and halting 7-Mode systems	3
Cabling the 7-Mode disk shelves	3
Importing 7-Mode data and configurations	4
Preproduction testing	4
Starting production	4
Committing the project	4

Copy-free transition workflow

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



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

Phases of copy-free transition

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

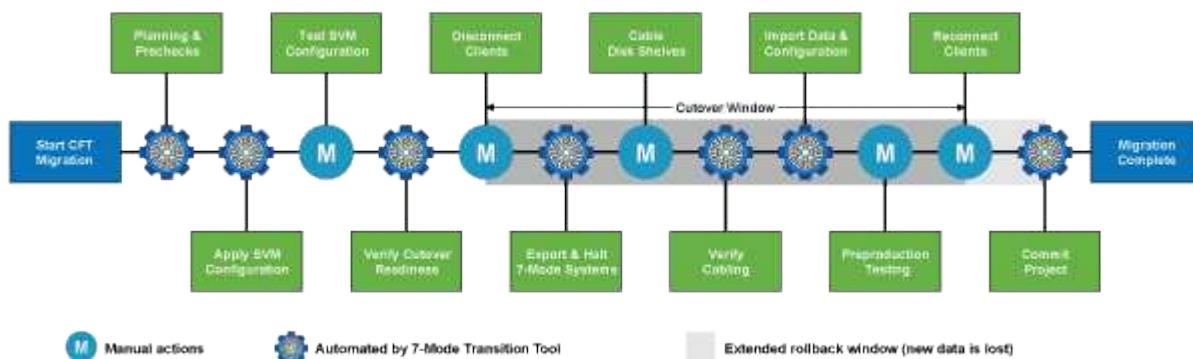
Copy-free transition is a disruptive operation. Therefore, you must plan for the downtime for the applications and workload running on the 7-Mode storage systems.

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



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

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



Planning the project

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

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



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

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

You must resolve any errors before continuing with the transition. Although the tool allows you to continue without resolving warnings, it is a best practice to address any warnings before continuing with the transition. You can run the prechecks multiple times to verify that all the errors have been resolved.

SVM provisioning

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

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

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

Exporting storage configurations and halting 7-Mode systems

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

The tool performs the following operations in the export phase:

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

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

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

Cabling the 7-Mode disk shelves

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



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

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

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

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

Importing 7-Mode data and configurations

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

The tool performs the following operations in the import phase:

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

Preproduction testing

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

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

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

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

Starting production

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

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

Committing the project

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

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

Related information

[NetApp Downloads: Config Advisor](#)

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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