



Transitioning Linux host file systems on LVM devices

ONTAP 7-Mode Transition

NetApp
October 09, 2025

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

Table of Contents

- Transitioning Linux host file systems on LVM devices 1
 - Testing LUNs with file systems on LVM devices before the cutover phase of copy-based transitions 1
 - Preparing for cutover phase when transitioning Linux host file systems on LVM devices 3
 - Mounting logical volumes on Linux hosts after transition 4

Transitioning Linux host file systems on LVM devices

When you transition a Linux host file system on a Logical Volume Manager (LVM), you must perform specific steps to prepare for the cutover phase and you must mount the logical volumes after the transition.

Related information

[Preparing for cutover phase when transitioning Linux host file systems on LVM devices](#)

[Mounting logical volumes on Linux hosts after transition](#)

Testing LUNs with file systems on LVM devices before the cutover phase of copy-based transitions

If you are using the 7-Mode Transition Tool (7MTT) 2.2 or later and Data ONTAP 8.3.2 or later to perform a copy-based transition of your Red Hat Enterprise Linux (RHEL) host, you can test your transitioned clustered Data ONTAP LUNs with file systems on LVM devices before the cutover phase. Your source host can continue to run I/O to your source 7-Mode LUNs during testing.

- Your new clustered Data ONTAP LUNs must be mapped to the test host.
- Your LUNs must be ready for transition.

You should maintain hardware parity between the test host and the source host and you should perform the following steps on the test host.

Your clustered Data ONTAP LUNs are in read/write mode during testing. They convert to read-only mode when testing is complete and you are preparing for the cutover phase.

During test mode you do not deactivate or export the volume group. For this reason, you might see file system errors when you mount the logical volumes on the test host.

Steps

1. After the baseline data copy is complete, select **Test Mode** in the 7MTT user interface (UI).
2. In the 7MTT UI, click **Apply Configuration**.
3. On the test host, discover your new clustered Data ONTAP LUNs:

```
rescan-scsi-bus.sh
```

4. Verify that your new clustered Data ONTAP LUNs have been discovered:

```
sanlun lun show
```

5. Configure DMMP devices for your clustered Data ONTAP LUNs:

```
multipath
```

6. Obtain the device handle ID for the clustered Data ONTAP LUNs:

```
multipath -ll
```

The following is an example of a device handle ID: "3600a09804d532d79565d47617679764d"

7. Identify the DMMP devices used by the LVM:

```
pvscan
```

3600a09804d532d79565d476176797655 is an example of a DMMP device used by the LVM.

8. Identify the volume group:

```
vgscan
```

9. Identify the logical volume:

```
lvscan
```

10. Enable the logical volumes: * **vgchange -ay volume_group**

11. Verify the logical volume status: * **lvdisplay**

The LV Status column in the output should display available.

12. Determine whether a mount point entry for the logical volume exists in the `/etc/fstab` file on the source host.

In the following example, logical volume `/dev/mapper/vg_7MTT-lv1` is displayed in the `/etc/fstab` file:

```
# /etc/fstab
...
tmpfs    /dev/shm tmpfs    defaults      0 0
devpts   /dev/pts devpts   gid=5, mode=620 0 0
sysfs    /sys     sysfs    defaults      0 0
proc     /proc    proc     defaults      0 0
/dev/mapper/vg_7MTT-lv1 /7MTT    ext4     defaults 0 0
```

13. If a mount point entry for the logical volume exists in the `/etc/fstab` file on the source host, manually edit the `/etc/fstab` file on the test host to add the mount point entry.

14. Mount the mount point:

```
mount -a
```

15. Verify that the mount points are mounted:

```
mount
```

16. Perform your testing as needed.

17. After you have completed your testing, shut down your host:

```
shutdown -h -t0 now
```

18. In the 7MTT UI, click **Finish Testing**.

If your clustered Data ONTAP LUNs are to be remapped to your source host, you must prepare your source host for the cutover phase. If your clustered Data ONTAP LUNs are to remain mapped to your test host, no further steps are required on the test host.

Related information

[Gathering pretransition information from the Inventory Assessment Workbook](#)

[Preparing for cutover phase when transitioning Linux host file systems on LVM devices](#)

Preparing for cutover phase when transitioning Linux host file systems on LVM devices

If you are transitioning a Linux host file system on a Logical Volume Manager (LVM) device, there are steps you must perform before the cutover phase.

- For FC configurations, you must have fabric connectivity and zoning to clustered Data ONTAP controllers.
- For iSCSI configurations, your iSCSI sessions must be discovered and logged in to your clustered Data ONTAP controllers.
- You must have the following pretransition information gathered from the *Inventory Assessment Workbook*:
 - The DMMP device names used by the LVM
 - The volume group name
 - The logical volume name
 - The file system configured on the logical volume device
 - The directory on which the logical volumes are mounted
- For copy-based transitions, perform these steps before initiating the Storage Cutover operation in the 7-Mode Transition Tool (7MTT).
- For copy-free transitions, perform these steps before initiating the Export & Halt 7-Mode operation in the 7MTT.

Steps

1. Stop I/O to LV mount points.
2. Shut down the applications accessing the LUNs according to application vendor's recommendations.
3. Unmount the LV mount point:

```
umount dir_name
```

4. Disable the logical volume:

```
vgchange -an vg_name
```

5. Verify the logical volume status:

```
lvdisplay dir_name
```

The LV status should display “NOT available”.

6. Export the volume group:

```
vgexport vg_name
```

7. Verify the VG status:

```
vgdisplay vg_name
```

The VG status should display “exported”.

8. Flush the 7-Mode DDMP device IDs:

```
multipath -f device_name
```

Related information

[Gathering pretransition information from the Inventory Assessment Workbook](#)

Mounting logical volumes on Linux hosts after transition

After the transition from ONTAP operating in 7-Mode to clustered Data ONTAP, your logical volumes are offline. You must mount those logical volumes for your LUNs to be accessible to your hosts.

If you are doing a copy-free transition (CFT), procedures for vol rehost must be complete. See the [7-Mode Transition Tool Copy-Free Transition Guide](#) for details.

- For copy-based transitions (CBTs), perform these steps after completing the Storage Cutover operation in the 7-Mode Transition Tool (7MTT).
- For CFTs, perform these steps after the Import Data & Configuration operation in the 7MTT.

1. Generate the 7-Mode to clustered Data ONTAP LUN mapping file:

- For copy-based transitions, run the following command from the Linux host where the 7MTT is installed:

```
transition cbt export lunmap -p project-name -o file_path
```

For example:

```
transition cbt export lunmap -p SanWorkLoad -o c:/Libraires/Documents/7-to-C-LUN-MAPPING.csv
```

- For copy-free transitions, run the following command from the system where the 7MTT is installed:

```
transition cft export lunmap -p p_roject-name_ -s svm-name -o output-file
```

For example:

```
transition cft export lunmap -p SanWorkLoad -s svml -0  
c:/Libraries/Documents/7-to-C-LUN-MAPPING-svml.csv
```



You must run this command for each of your storage virtual machines (SVMs).

2. Remove the SCSI devices created for 7-Mode LUNs:

- To remove all of the SCSI devices:

```
rescan-scsi-bus.sh -r
```

- To remove each SCSI device individually:

```
echo 1> /sys/block/SCSI_ID/delete
```

This command must be executed on all 7-Mode LUN SCSI devices. See the SCSI Device ID column on the SAN Host LUNs tab of the *Inventory Assessment Workbook* to identify the SCSI device IDs for the LUNs.

3. Discover new ONTAP LUNs:

```
rescan-scsi-bus.sh
```

4. Configure DMMP devices for ONTAP LUNs:

```
multipath
```

5. Verify that ONTAP LUNs are discovered:

```
sanlun lun show
```

6. Determine the new ONTAP LUN device handle ID:

```
multipath -ll Device_handle_name
```

7. Import the volume group:

```
vgimport vg_name
```

8. Verify the volume group status:

```
vgdisplay
```

9. Enable logical volumes:

```
vgchange -ay vg_name
```

10. Verify the logical volume status:

```
lvdisplay
```

The LV status should be displayed as “available”.

11. Mount the logical volumes from the ONTAP LUN to its respective mount point directory:

```
mount lv_namemount_point
```

If the mount points are defined in the `etc/fstab` file, you can use the `mount -a` command to mount the logical volumes.

12. Verify the mount points:

mount

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.