



Transitioning HP-UX host LUNs with file systems

ONTAP 7-Mode Transition

NetApp
May 31, 2021

This PDF was generated from https://docs.netapp.com/us-en/ontap-7mode-transition/san-host/task_preparing_to_transition_hp_ux_host_luns_with_file_systems.html on May 31, 2021. Always check docs.netapp.com for the latest.

Table of Contents

- Transitioning HP-UX host LUNs with file systems 1
 - Preparing to transition HP-UX host LUNs with file systems 1
 - Testing data LUNs on HP-UX hosts before the cutover phase of copy-based transitions 3
 - Preparing for cutover phase when transitioning HP-UX host data LUNs with file systems 4
 - Mounting HP-UX host LUNs with file systems after transition 5

Transitioning HP-UX host LUNs with file systems

If you transition an HP-UX host LUN with a file system from Data ONTAP operating in 7-Mode to clustered Data ONTAP using the 7-Mode Transition Tool (7MTT), you must perform specific steps before and after transition to remediate transition issues on the host.

Preparing to transition HP-UX host LUNs with file systems

Before you transition HP-UX host LUNs with file systems from Data ONTAP operating in 7-Mode to clustered Data ONTAP, you must gather information you need for the transition process.

Steps

1. Display your LUNs to identify the name of the LUNs to be transitioned:

```
lun show
```

2. Locate the SCSI device name for the LUNs to be transitioned and the Agile name for the SCSI device:

```
sanlun lun show -p
```

In the following example, the transition LUNs are lun1 and lun3. The SCSI device names for lun1 are `/dev/dsk/c14t0d1`, `/dev/dsk/c27t0d1`, `/dev/dsk/c40t0d1`, and `/dev/dsk/c31t0d1`. The SCSI device names for lun3 are `/dev/dsk/c14t0d2`, `/dev/dsk/c27t0d2`, `/dev/dsk/c40t0d2`, and `/dev/dsk/c31t0d2`.

The Agile name for SCSI device `/dev/dsk/c31t0d1` is `/dev/rdisk/disk11`.

```
ONTAP Path: f8040-211-185:/vol/vol1185_n1/lun3
  LUN: 1
  LUN Size: 3g
  Host Device: /dev/rdisk/disk11
  Mode: 7
Multipath Provider: None
host      vservers  /dev/dsk
path      path      filename      host      vservers
state     type      or hardware path  adapter  LIF
-----  -
up        secondary /dev/dsk/c14t0d1  fcd0      fc4
up        primary   /dev/dsk/c27t0d1  fcd0      fc2
up        primary   /dev/dsk/c40t0d1  fcd1      fc1
up        secondary /dev/dsk/c31t0d1  fcd1      fc3
```

```

ONTAP Path: f8040-211-183:/vol/vol183_n1/lun1
      LUN: 3
      LUN Size: 3g
Host Device: /dev/rdisk/disk14
      Mode: 7
Multipath Provider: None
host      vserver    /dev/dsk
path      path        filename      host      vserver
state     type        or hardware  path      adapter   LIF
-----
up        secondary /dev/dsk/c14t0d1 fcd0     fc4
up        primary   /dev/dsk/c27t0d1 fcd0     fc2
up        primary   /dev/dsk/c40t0d1 fcd1     fc1
up        secondary /dev/dsk/c31t0d1 fcd1     fc3

```

3. Identify the WWID for the LUN on the host:

```
scsimgr get_info -D Agile_name_for_SCSI_device
```

In this example, the LUN WWID for device `/dev/rdisk/disk11` is `0x600a09804d537739422445386b755529`:

```

bash-2.05# scsimgr get_info -D /dev/rdisk/disk11 |grep WWID
World Wide Identifier (WWID)      = 0x600a09804d537739422445386b755529

```

4. List and record your volume groups:

```
vgdisplay
```

5. List and record your volume groups, logical volumes, and physical volumes:

```
vgdisplay -v vg_name
```

6. Write the VGID and logical volumes for the volume group to a mapfile:

```
vgexport -p -s -m /tmp/mapfile/vg01 vg01
```

7. Make a backup copy of the `mapfile.vg01` to an external source.

8. List and record the mount points:

```
bdf
```

The following example shows how the mount points should be displayed:

```
bash-2.05# bdf
Filesystem      kbytes      used        avail      used  Mounted on
/dev/vg01/lvol1 123592960   1050952    22189796   5%   /mnt/qa/vg01
/dev/vg01/lvol2 23592960    588480     22645044   3%   /mnt/qa/vg02
```

Testing data LUNs on HP-UX hosts 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 HP-UX host data LUNs, you can test your transitioned clustered Data ONTAP LUNs to verify that you can mount your MPIO device before the cutover phase. Your source host can continue to run I/O to your source 7-Mode LUNs during testing.

Your new ONTAP LUNs must be mapped to your test host and 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 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.

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, rescan your new ONTAP LUNs:

```
ioscan -fnC disk
```

4. Verify that your ONTAP LUNs are present:

```
sanlun lun show
```

5. Copy the `/tmp/mapfile.vg01 mapfile` previously copied to your external source to your new host.
6. Use the mapfile to import the volume group:

```
vgimport -s -m /tmp/mapfile/vg01 vg01
```

7. Verify that the `VG Status` is displayed as `available`:

```
vgdisplay
```

8. Convert the legacy Device Special Filename (DSF) to persistent DSF:

```
vgdsf -c /dev/vg01
```

9. Use the mount command to manually mount each of the logical volumes.

10. Run the `fsck` command if you are prompted to do so.

11. Verify the mount points:

```
bdf
```

12. Perform your testing as needed.

13. Shut down the test host.

14. In the 7MTT UI, click **Finish Test**.

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

Preparing for cutover phase when transitioning HP-UX host data LUNs with file systems

If you are transitioning an HP host data LUN with a file system from Data ONTAP operating in 7-Mode to clustered Data ONTAP, you must perform certain steps before entering the cutover phase.

If you are using an FC configuration, fabric connectivity and zoning to the clustered Data ONTAP nodes must be established.

If you are using an iSCSI configuration, the iSCSI sessions to the clustered Data ONTAP nodes must be discovered and logged in.

For copy-based transitions, perform these steps after completing the Storage Cutover operation in the 7-Mode Transition Tool (7MTT). Copy-free transitions are not supported for HP-UX hosts.

Steps

1. Stop I/O on all mount points.
2. Shut down each application accessing the LUNs according to the recommendations of the application vendor.
3. Unmount all of the mount points:

```
umount mount_point
```

4. Export your volume group and write the VGID and logical volumes for the volume group to a mapfile:

```
vgexport -p -s -m /tmp/mapfile.vg01 vg01
```

5. Make a backup copy of the mapfile.vg01 file to an external source.
6. Disable the volume group:

```
vgchange -a n vg_name
```

7. Export the volume group:

```
vgexport vg_name
```

8. Verify that the volume group has been exported:

```
vgdisplay
```

The exported volume group information should not be displayed in the output.

Mounting HP-UX host LUNs with file systems after transition

After transitioning HP-UX host LUNs with file systems from Data ONTAP operating in 7-Mode to clustered Data ONTAP, you must mount the LUNs.

For copy-based transitions, perform these steps after completing the Storage Cutover operation in the 7-Mode Transition Tool (7MTT). Copy-free transitions are not supported for HP-UX hosts.

Steps

1. Discover new clustered Data ONTAP LUNs:

```
ioscan -fnC disk
```

2. Verify that the clustered Data ONTAP LUNs have been discovered:

```
sanlun lun show
```

3. Verify that the `lun-pathname` for the clustered Data ONTAP LUNs is the same as the `lun-pathname` for the 7-Mode LUNs prior to transition.

4. Verify that the output in the mode column has changed from `7` to `C`.

5. Use the `mapfile` file to import the volume group:

```
vgimport -s -v -m /tmp/mapfile.vg01 /dev/vg01"
```

6. Activate the logical volumes:

```
vgchange -a y vg_name
```

7. Convert the legacy Device Special Filename (DSF) to persistent DSF:

```
vgdsf -c /dev/vg01
```

8. Verify that the VG Status is displayed as available:

```
vgdisplay
```

9. Manually mount each of the devices:

```
mount -F vxfs -o largefiles device_name mount_point
```

10. Run the `fsck` command if you are prompted to do so.

11. Verify the mount points:

bdf

The following example shows how the mount points should be displayed:

```
bash-2.05# bdf
Filesystem      kbytes    used    avail    used  Mounted on
/dev/vg01/lvol1 23592960 1050952 22189796 5%    /mnt/qa/vg01
/dev/vg01/lvol2 23592960  588480 22645044 3%    /mnt/qa/vg02
```


Copyright Information

Copyright © 2021 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, 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.