



Transitioning of SAN boot LUNs

ONTAP 7-Mode Transition

NetApp
May 31, 2021

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

Table of Contents

- Transitioning of SAN boot LUNs 1
 - Types of SAN boot LUNs supported for transition 1
 - Preparing for transition of FC or FCoE SAN boot LUNs on RHEL hosts 1
 - Preparing for transition of iSCSI SAN boot LUNs 3
 - Testing SAN boot LUNs on RHEL hosts before the cutover phase of copy-based transitions 5
 - Preparing for the cutover phase when transitioning SAN boot LUNs 6
 - Discovering SAN boot LUNs after transition 6

Transitioning of SAN boot LUNs

You must reboot SAN boot LUNs before you transition from Data ONTAP operating in 7-Mode to clustered Data ONTAP using the 7-Mode Transition Tool (7MTT). You must perform specific steps to prepare for the cutover phase, and after transition you must discover your LUNs.

Related information

[Preparing for transition of FC or FCoE SAN boot LUNs on RHEL hosts](#)

[Preparing for transition of iSCSI SAN boot LUNs](#)

[Discovering SAN boot LUNs after transition](#)

Types of SAN boot LUNs supported for transition

Only certain types of SAN boot LUNs are supported for transition from Data ONTAP operating in 7-Mode to clustered Data ONTAP.

The following SAN boot LUNs are supported for transition:

- FC or FCoE SAN boot LUNs
- iSCSI SAN boot LUNs for Red Hat Enterprise Linux (RHEL) 6

Transition of iSCSI SAN boot LUNs for RHEL 5.x is not supported.

Preparing for transition of FC or FCoE SAN boot LUNs on RHEL hosts

Before you transition an FC or FCoE SAN boot LUN, you must perform specific steps on your Red Hat Enterprise Linux (RHEL) host.

You must have the following information from the *Inventory Assessment Workbook*:

- 7-Mode LUN name on which RHEL 5 or RHEL 6 is installed
- SCSI device name for the transition LUN
- DMMP device name for the transition LUN
- Mount directory
- File system configured on the DMMP device
- UUID number of the /boot partition
- Name of the `initrid` image

This procedure applies to copy-based transitions and copy-free transitions.

1. Verify that the DMMP device exists in the /dev/mapper directory:

```
ls /dev/mapper/ DMMP_device_name
```

If you cannot locate the DMMP device, then it might be using an alias or user-friendly name.

2. Identify the DMMP devices and Logical Volume Manager (LVM) names on which the RHEL 5 or RHEL 6 operating system /boot and root (/) directories are installed:

```
df - h
```

By default, RHEL 5 and RHEL 6 are installed on the root (/) partition on the logical volume. If the root partition is installed on the logical volume, then no pretransition changes to the configuration are required.

3. If the /boot partition is installed on the DMMP device, confirm how the /boot partition is referenced to mount in `/etc/fstab` at boot time.
4. If the /boot partition is referenced in `/etc/fstab` by its DMMP device name, replace the DMMP device name with the file system UUID name.
5. Make a backup of the `/etc/fstab` file:

```
cp /etc/fstab /etc/fstab_pre_transition_file_name
```

6. Edit the `/etc/fstab` file to replace the DMMP device name with its respective file system UUID number.
7. Make a backup of the `initrd` image file:

```
cp /boot/initrd_image_file_nameinitrd_image_file_name.bak
```

8. For RHEL 5 only:

- a. In the `/etc/multipath.conf` file, identify the SWAP partition device.

In the following example, `/dev/VolGroup00/LogVol01` is the SWAP partition device:

```
/dev/VolGroup00/LogVol01 swap swap defaults 0 0
```

- b. Create a label for mounting the swap partition:

```
swapoff swap-partition_device
```

```
mkswap -L label-for-swapswap-partition-device
```

```
swapon swap-partition_device
```

- c. Replace the SWAP partition device name in the `/etc/fstab` file with THE SWAP label.

The updated line in the `/etc/fstab` file should be as follows:

```
LABEL=SwapPartition swap swap defaults 0 0
```

9. Re-create the `initrd` image.

- For RHEL5:

```
mkinitrd -f/boot/ initrd-"`uname-r`".img `uname-r` --with multipath
```

- For RHEL 6:
`dracut --force --add multipath --verbose`

10. Restart the host to boot from the new `initrd` image.

Related information

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

Preparing for transition of iSCSI SAN boot LUNs

Before you transition an iSCSI SAN boot LUN, you must perform specific steps on the host. Transition of Red Hat Enterprise Linux (RHEL) 5.x is not supported. Transition of RHEL 6 is supported.

You must have the following information from the *Inventory Assessment Workbook*:

- Name of the LUN that RHEL 6 is installed on
- DMMP device name for the transition LUN
- Logical volume (LV) name
- Volume group (VG) name
- Physical volume (PV) devices
- Logical Volume Manager (LVM) names and mount directories on which RHEL 6 /boot and root (/) partitions are installed
- File system configured on DMMP
- iSCSI sessions for 7-Mode controllers
- Grub information
- IQN number of the storage virtual machine (SVM) where the iSCSI SAN boot LUN will be created
- LIF IP address of the clustered Data ONTAP SVM where the iSCSI SAN boot LUN will be created

This procedure applies to copy-based transitions and copy-free transitions.

Steps

1. Verify that the DMMP device exists in the `/dev/mapper` directory:

```
ls /dev/mapper/DMMP_device_name
```

If the DMMP device is not displayed, the device might be using an alias or a user-friendly name.

2. Determine whether the DMMP device is part of an LVM:

```
blkid
```

If the DMMP device `TYPE` value is `LVM2_member`, the DMMP is part of an LVM.

3. Obtain the mount point details of the `/` and `/boot` partitions from the `/etc/fstab` file:

- If the `/boot` partition is installed on a DMMP device, check how it is referenced to mount in the `/etc/fstab` file at boot time.

- If the `/boot` partition is mounted using the file system UUID that you obtained using the `blkid` command output, then no pretransition change are required.
- 4. If the `/boot` partition is referenced in `/etc/fstab` file by its DMMP device name, replace the DMMP device name with the file system UUID name.
- 5. For iSCSI SAN booted hosts, edit the `/boot/grub/grub.conf` file to create a new kernel command-line entry that includes the clustered Data ONTAP controller's IQN number and iSCSI session information.

This example shows the `/boot/grub/grub.conf` file before editing. The kernel command line has the 7-Mode controller's IQN number and iSCSI session information.

```
title Red Hat Enterprise Linux Server (2.6.32-431.el6.x86_64)
    root (hd0,0)
    kernel /vmlinuz-2.6.32-431.el6.x86_64 ro
root=/dev/mapper/vg_ibmx3550m3229-LogVol100 ifname=eth0:5c:f3:fc:ba:46:d8
rd_NO_LUKS netroot=iscsi:@10.226.228.241::3260::iqn.1992-
08.com.netapp:sn.1574168453 LANG=en_US.UTF-8
rd_LVM_LV=vg_ibmx3550m3229/LogVol01 rd_LVM_LV=vg_ibmx3550m3229/LogVol100
rd_NO_MD netroot=iscsi:@10.226.228.155::3260::iqn.1992-
08.com.netapp:sn.1574168453 iscsi_initiator= iqn.1994-
08.com.redhat:229.167 crashkernel=auto ip=eth0:dhcp
    initrd /initramfs-2.6.32-431.el6.x86_64.img
```

This example shows the `/boot/grub/grub.conf` file after adding a new title with the cDOT suffix, and the new kernel command line with the clustered Data ONTAP controller's IQN number and iSCSI session information:

```
title Red Hat Enterprise Linux Server (2.6.32-431.el6.x86_64) - cDOT
    root (hd0,0)
    kernel /vmlinuz-2.6.32-431.el6.x86_64 ro
root=/dev/mapper/vg_ibmx3550m3229-LogVol100 ifname=eth0:5c:f3:fc:ba:46:d8
rd_NO_LUKS netroot=iscsi:@10.226.228.99::3260:: ::iqn.1992-
08.com.netapp:sn.81c4f5cc4aa611e5b1ad00a0985d4dbe:vs.15 LANG=en_US.UTF-8
rd_LVM_LV=vg_ibmx3550m3229/LogVol01 rd_LVM_LV=vg_ibmx3550m3229/LogVol100
rd_NO_MD netroot=iscsi:@10.226.228.98::3260:: ::iqn.1992-
08.com.netapp:sn.81c4f5cc4aa611e5b1ad00a0985d4dbe:vs.15
netroot=iscsi:@10.226.228.97::3260:: ::iqn.1992-
08.com.netapp:sn.81c4f5cc4aa611e5b1ad00a0985d4dbe:vs.15
netroot=iscsi:@10.226.228.96::3260:: ::iqn.1992-
08.com.netapp:sn.81c4f5cc4aa611e5b1ad00a0985d4dbe:vs.15 iscsi_initiator=
iqn.1994-08.com.redhat:229.167 crashkernel=auto ip=eth0:dhcp
    initrd /initramfs-2.6.32-431.el6.x86_64.img
```

- 6. Back up the existing `initramfs` file.

```
# cd /boot
# cp initramfs-2.6.32-71.el6.x86_64.img initramfs-2.6.32-
71.el6.x86_64.img.img_bak
```

7. Update the 7-Mode kernel line in the `/boot/grub/grub.conf` file with the backup `initrd` image name.

For RHEL 6.4 and later, verify that the clustered Data ONTAP kernel line is appended with “`rdloaddriver=scsi_dh_alua`” in the `/boot/grub/grub.conf` file.

8. If the `/boot/grub/grub.conf` file is updated, then update the kernel initial RAM disk (`initramfs`).

The `initramfs` file must be re-created so that the new clustered Data ONTAP IQN number and iSCSI sessions are referenced, and so that the host establishes an iSCSI connection with clustered Data ONTAP controllers at boot time.

9. Re-create the `initrd` image by using the `dracut -force --add multipath --verbose` command.

Related information

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

Testing SAN boot LUNs on RHEL 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 your Red Hat Enterprise Linux (RHEL) host, you can test your transitioned ONTAP SAN boot LUNs 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.

- For copy-based transitions, you must perform these steps after completing the Storage Cutover operation in the 7-Mode Transition Tool.
- For copy-free transitions, you must perform these steps after the Import Data and Configuration operation in the 7-Mode Transition Tool.

Steps

1. For FC and FCoE configurations only:
 - a. Enter the HBA BIOS settings mode.
 - b. Choose **Rescan** to discover the ONTAP SAN boot LUNs on the host.
 - c. Remove the 7-Mode boot LUN ID.
 - d. Add the ONTAP boot LUN ID in the HBA BIOS.
 - e. Exit the HBA BIOS settings mode, and then reboot the host.
2. After the host reboots, change the IP address and host name on the test host.

3. Verify that your new ONTAP LUNs have been discovered:

```
sanlun lun show
```

4. Configure DMMP devices for your ONTAP LUNs:

```
multipath -ll
```

5. Perform your testing as needed.

6. Shut down the test host:

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

7. In the 7-Mode Transition Tool user interface (UI), click **Finish Testing**.

If your ONTAP LUNs are to be remapped to your source host, you must prepare your source host for the cutover phase. If your 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 the cutover phase when transitioning SAN boot LUNs](#)

Preparing for the cutover phase when transitioning SAN boot LUNs

If you are transitioning SAN boot LUNs from Data ONTAP operating in 7-Mode to clustered Data ONTAP, there are certain prerequisites you must be aware of before entering the cutover phase.

You must have fabric connectivity and zoning to your clustered Data ONTAP controllers for FC configurations. For iSCSI configurations, your iSCSI sessions must be discovered and logged in to your clustered Data ONTAP controllers. You must also shut down your host.

- For copy-based transitions, you should shut down your host before initiating the Storage Cutover operation in the 7-Mode Transition Tool (7MTT). Copy-free transitions are not supported on HP-UX hosts.
- For copy-free transitions, you should shut down your host before initiating the Export & Halt 7-Mode operation in the 7MTT.

Discovering SAN boot LUNs after transition

After transition of your SAN boot LUNs from Data ONTAP operating in 7-Mode to clustered Data ONTAP, you must discover the SAN boot LUNs on your host. This is necessary for copy-based transitions (CBTs) and copy-free transitions (CFTs). This applies to FC, FCoE, and iSCSI configurations.

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

1. Boot the host.
2. For FC and FCoE configurations only:
 - a. Enter the HBA BIOS settings mode.
 - b. Choose **Rescan** to discover the clustered Data ONTAP SAN boot LUNs on the host.
 - c. Remove the 7-Mode boot LUN ID.
 - d. Add the clustered Data ONTAP boot LUN ID in the HBA BIOS.
 - e. Exit the HBA BIOS settings mode, and then reboot the host.
3. After the reboot is complete, verify the clustered Data ONTAP LUNs:

```
sanlun lun show
```

4. Verify the DMMP device:

```
multipath -ll
```

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.