



RHEL host remediation

ONTAP 7-Mode Transition

NetApp

February 11, 2024

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

Table of Contents

- RHEL host remediation 1
 - Gathering pretransition information from the Inventory Assessment Workbook 1
 - Transitioning RHEL DMMP devices without file systems 2
 - Transitioning LUNs with mount points using DMMP device names 7
 - Transitioning LUNs with mount points using DMMP alias names 14
 - Transitioning Linux host file systems on LVM devices 20
 - Transitioning of SAN boot LUNs 25

RHEL host remediation

If you are using the 7-Mode Transition Tool (7MTT) to move from Data ONTAP operating in 7-Mode to clustered Data ONTAP in a SAN environment, you must perform a series of steps on your Red Hat Enterprise Linux (RHEL) hosts based on your LUN type before and after the transition to avoid transition complications.

7MTT transitions only support RHEL 5 and RHEL 6.

Related information

[Transitioning RHEL DMMP devices without file systems](#)

[Transitioning LUNs with mount points using DMMP device names](#)

[Transitioning LUNs with mount points using DMMP alias names](#)

[Transitioning Linux host file systems on LVM devices](#)

[Transitioning of SAN boot LUNs](#)

Gathering pretransition information from the Inventory Assessment Workbook

There is information in the *Inventory Assessment Workbook* that you need in various phases of the transition. You should gather and record this information before you begin the transition so that you can reference it as needed throughout the process.

Steps

1. Use the Inventory Collect Tool (ICT) to generate the *Inventory Assessment Workbook*.
2. Open the *Inventory Assessment Workbook*.
3. Go to the **LUN** tab.
4. In the **LUN name** column, identify and record the name of the LUN to be transitioned.
5. Go to the **SAN Host LUNs** tab.
6. In the **SCSI Device ID** column, identify and record the SCSI device name.
7. In the **OS Device ID** column, identify and record the DMMP device name for the LUN to be transitioned.
8. In the **Filesystems** column, identify and record the file system configured on the DMMP device.
9. In the **UUID** column, identify and record the UUID number for the LUN.
10. In the **Mount** column, identify and record the directory on which the DMMP device is mounted.
11. Go to the **LVMs** tab.
12. In the **Physical Volume Name** column, identify and record the DMMP devices that are used by the logical volume.
13. Go to the **SAN Host LVMs** tab.
14. In the **Volume Group Name** column, identify and record the volume group.

15. In the **Logical Volume Path** column, identify and record the logical volume.
16. Go to the **SAN Host Filesystems** tab.
17. In the **Filesystem** column, identify and record the file system configured on the logical volume.
18. In the **Mount** column, identify and record the directory to which the logical volumes are mounted.
19. Go to the **GRUB Configuration** tab.
20. In the **initrd** column, identify and record the initrd image to be modified.
21. Go to the **SAN Host HBAs** tab.

You can also see the **iSCSI SAN Interfaces** tab to identify the iSCSI IQN number and IP addresses that are configured on 7-Mode controllers.

22. In the **Target IPs (iSCSI)** column identify and record the iSCSI sessions for the 7-Mode controllers.

Related information

[What the Inventory Collect Tool is](#)

What the Inventory Collect Tool is

The Inventory Collect Tool (ICT) is a stand-alone utility for collecting configuration and inventory information about 7-Mode storage controllers, hosts attached to controllers, and applications running on these hosts for assessing the transition readiness of these systems. You can use the ICT to generate information about your LUNs and the configuration that you need for transition.

The ICT generates an *Inventory Assessment Workbook* and an Inventory Report XML file that contains configuration details of the storage and host systems.

The ICT is available for ESXi, 5.x, ESXi 6.x, and Windows hosts.

Transitioning RHEL DMMP devices without file systems

Before transition of a Red Hat Enterprise Linux (RHEL) DMMP device without a file system, you must verify that the DMMP device does not have a file system. You must also perform specific steps to prepare for the cutover phase, and after transition you must replace the WWID.

Related information

[Verifying that RHEL LUNs are ready for transition using the Inventory Assessment Workbook](#)

[Verifying that RHEL 5 LUNs are ready for transition using the CLI](#)

[Verifying that RHEL 6 DDMP devices are ready for transition using CLI](#)

[Preparing for cutover when transitioning a Linux host DMMP device without a file system](#)

[Replacing 7-Mode LUN WWIDs on Linux hosts after transition of LUNs](#)

Verifying that RHEL LUNs are ready for transition using the Inventory Assessment Workbook

If your Red Hat Enterprise Linux (RHEL) 5 or RHEL 6 LUN is configured with a device mapper multipath (DMMP), you should verify that a file system is not configured before you transition the LUN from Data ONTAP operating in 7-Mode to clustered Data ONTAP.

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

Steps

1. Gather pretransition information from the *Inventory Assessment Workbook*.
2. Check whether the DMMP device entry is present under the **SAN Host File system** tab.

If the DMMP device entry is not present, a file system is not configured and you can transition the LUN.

Verifying that RHEL 5 LUNs are ready for transition using the CLI

If your Red Hat Enterprise Linux (RHEL) 5 LUN is configured with a device mapper multipath (DMMP), you should verify that a file system is not configured before you transition the LUN from Data ONTAP operating in 7-Mode to clustered Data ONTAP.

Steps

1. Locate the SCSI device name for the LUN to be transitioned:

```
sanlun lunshow
```

2. Identify the DMMP device name for the LUN:

```
multipath -ll
```

The DMMP device name can be a device handle ID (WWID) such as 360a980003753456258244538554b4b53, or it can be an alias, such as, dmmp_raw_lun.

3. Verify that the LUN does not have a file system:

```
dumpe2fs/dev/mapper/DMMP device name
```

If the LUN does not have a file system, Couldn't find valid filesystem superblock is displayed in the output.

Verifying that RHEL 6 DDMP devices are ready for transition using CLI

Before you transition your Red Hat Enterprise Linux (RHEL) 6 DMMP device, you must verify that it is not part of a Logical Volume Manager (LVM) and that it does not have a file system.

Steps

1. Gather pretransition information from the *Inventory Assessment Workbook*.
2. 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.

3. Determine whether the DMMP device is part of an LVM and whether the DMMP device has a file system:

```
blkid
```

If the DMMP device is not part of an LVM and has no file system the device entry should not be displayed in blkidoutput.

Testing DMMP devices without file systems 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 transition your Red Hat Enterprise Linux (RHEL) 5 host, you can test your transitioned clustered Data ONTAP LUNs to verify that you can bring your host and applications online 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 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 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.

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. Rescan your new clustered Data ONTAP LUNs on the test host:

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

4. Obtain the new SCSI device names for the clustered Data ONTAP LUNs:

```
sanlun lun show
```

In the following example, /dev/sdl is the SCSI device name for the lun_dmmp_raw LUN, and /dev/sdk is the SCSI device name for the lun_dmmp_raw_alias LUN:

```
[root@ibmx3550-229-108 /]# sanlun lun show
controller(7mode/E-Series)/
vserver (cDOT/FlashRay lun-pathname      filename
-----
vs_brb    /vol/dmmp_raw_vol/lun_dmmp_raw           /dev/sdl
vs_brb    /vol/dmmp_raw_alias_vol/lun_dmmp_raw_alias /dev/sdk
```

5. Obtain the device handle IDs (WWIDs) for the clustered Data ONTAP LUNs:

```
/sbin/scsi_id -g-u-s /block/SCSI_device_name
```

The following is an example of a WWID: "3600a09804d532d79565d47617679764d"

6. Check whether an alias is defined in the `/etc/multipath.conf` file on the source host.
7. If there is an alias defined on the source host, add the alias to the `/etc/multipath.conf` file on the test host, but replace the 7-Mode device handle ID with the clustered Data ONTAP LUN ID.
8. Update the DMMP alias settings:

```
multipath
```

9. Verify that the DMMP alias name correctly references the clustered Data ONTAP LUN:

```
multipath -ll
```

10. Perform your testing as needed.
11. After you have completed your testing, shut down the test host:

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

12. 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](#)

[Verifying that RHEL LUNs are ready for transition using the Inventory Assessment Workbook](#)

[Preparing for cutover when transitioning a Linux host DMMP device without a file system](#)

Preparing for cutover when transitioning a Linux host DMMP device without a file system

If you are transitioning a DMMP device without a file system from a Linux host, there are several steps you must perform before entering 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.

- 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 7MTT.

Steps

1. Stop I/O to the mount points.
2. Shut down the applications that are accessing the LUNs according to the application vendor's recommendations.
3. Flush the 7-Mode LUN DMMP device or alias:

```
multipath -f device_name
```

If needed, you can get the DMMP device name from the **OS Device ID** column under the SAN Host LUNs tab in the *Inventory Assessment Workbook*.

Replacing 7-Mode LUN WWIDs on Linux hosts after transition of LUNs

After LUN transition, the 7-Mode LUN WWID changes. You must replace it with the corresponding ONTAP LUN WWID before you can begin servicing data.

If you are doing a copy-free transition (CFT), then 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 7MTT.
- For CFTs, perform these steps after the Import Data & Configuration operation in the 7MTT is complete.

Steps

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

- For CBTs, 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 CFTs, run the following command from the system where the 7MTT is installed:

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

For example:

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



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

2. Make a note of the new ONTAP LUN device handle ID from the LUN mapping file.
3. 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.

4. Discover new ONTAP LUNs:

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

5. Identify the SCSI devices of the new ONTAP LUNs:

```
sanlun lun show
```

6. Get the WWIDs for the new ONTAP LUNs:

```
/lib/udev/scsi_id -g -u -d /dev SCSI_dev
```

7. If a DMMP alias is defined, then update the /etc/multipath.conf file to replace the 7-Mode LUN WWID with its corresponding ONTAP LUN WWID, so that the DMMP alias points to the clustered Data ONTAP LUN:

```
cat /etc/multipath.conf
```

8. Configure the DMMP devices:

```
multipath
```

9. Verify that the DMMP alias is correctly referencing the ONTAP LUN WWID:

```
multipath -ll
```

In the following sample output, the DMMP alias `dmmp_raw_lun` is referencing `3600a098051764b2d4f3f453135452d31` as the ONTAP WWID:

```
root@IBMX3550M3-229-169 ~]# multipath -ll dmmp_raw_lun
dmmp_raw_lun (3600a098051764b2d4f3f453135452d31) dm-8 NETAPP, LUN C-Mode
[size=1.0G] [features=3 queue_if_no_path pg_init_retries 50]
[hwhandler=1 alua] [rw]
\_round-robin 0 [prio=50][enabled]
  \_5:0:0:6 sdx 65:112 [active][ready]
    \_8:0:0:6 sdab 65:176 [active][ready]
\_round-robin 0 [prio=10][enabled]
  \_6:0:0:6 sdy 65:128 [active][ready]
    \_7:0:0:6 sdaa 65:160 [active][ready]
```

Transitioning LUNs with mount points using DMMP device names

Before transitioning a LUN with a mount point using a DMMP device name, you must replace the DMMP device name with its corresponding file system UUID number. You

must perform specific steps to prepare for the cutover phase and you must remount the DMMP devices on the host after transition. You perform the same procedures for Red Hat Enterprise Linux (RHEL) 5 and RHEL 6.

Related information

[Preparing RHEL LUNs with mount points using DMMP device names for transition using the Inventory Assessment Workbook](#)

[Preparing RHEL LUNs with mount points using DMMP alias names for transition using the CLI](#)

[Preparing for the cutover phase when transitioning LUNs with mount points using DMMP device names on Linux hosts](#)

[Remounting DMMP devices on Linux hosts after transition](#)

Preparing RHEL LUNs with mount points using DMMP device names for transition using the Inventory Assessment Workbook

Before transition of a LUN with a mount point using a DMMP device name, you must replace the DMMP device name with its respective file system UUID number. This applies to Red Hat Enterprise Linux (RHEL) 5 and RHEL 6.

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

Steps

1. Gather pretransition information from the *Inventory Assessment Workbook*.

Specifically, you need the following information:

- The file system configured on the DMMP device
- The directory on which the DMMP device is mounted
- The file system UUID for the DMMP device

Steps

1. Verify that the mount points for the DMMP device are defined in the `/etc/fstab` file.
2. Create a backup of the file:

```
cp /etc/fstab /etc/fstab_pre_transition
```

3. Edit the `/etc/fstab` file to replace the DMMP device name with its respective file system UUID number.

In the following example, the DMMP device `/dev/mapper/360a9800037534562572b453855496b41` is replaced by UUID `a073547e-00b6-4bf9-8e08-5eef08499a9c`:

```
[root@IBMX3550M3-229-169 ~]# cat /etc/fstab
/dev/VolGroup00/LogVol100 / ext3 defaults 1 1
LABEL=/boot /boot ext3 defaults 1 2
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/VolGroup00/LogVol101 swap swap defaults 0 0
/dev/mapper/test_vg-test_lv /mnt/lvm_ext3 ext3 defaults,_netdev 0 0
UUID=a073547e-00b6-4bf9-8e08-5eef08499a9c /mnt/dmmp_ext3 ext3
defaults,_netdev 0 0
```

Related information

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

Preparing RHEL LUNs with mount points using DMMP alias names for transition using the CLI

Before transition of a mount point using a DMMP device name, you must replace the DMMP device name with its respective file system UUID number.

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

Steps

1. Identify and record the SCSI device ID for the LUN to be transitioned:

```
sanlun lun show
```

The SCSI device ID is listed under the filename column in the output.

2. Identify and record the DMMP device name for the LUN to be transitioned:

```
multipath -ll SCSI_device_ID
```

In the following example, 360a9800037534562572b453855496b41 is the DMMP device name:

```
[root@IBMX3550M3-229-169 ~]# multipath -ll /dev/sdc
dmmp_fs_lun (360a9800037534562572b453855496b41) dm-3 NETAPP, LUN
[size=1.0G] [features=3 queue_if_no_path pg_init_retries 50]
[hwhandler=0][rw]
\_ round-robin 0 [prio=2][active]
  \_ 9:0:0:1 sdc 8:32 [active][ready]
  \_ 9:0:0:1 sdg 8:96 [active][ready]
```

3. Identify the file system configured on the DMMP device:

```
blkid | grep -i DMMP_device_name
```

The TYPE value in the output identifies the file system.

In the following example, the file system is ext3.

```
[root@ibmx3550-229-108 ~]#blkid | grep -i
3600a09804d532d79565d47617679658
/dev/mapper/3600a09804d532d79565d47617679658:
UUID="450b999a-4f51-4828-8139-29b20d2f8708" TYPE="ext3" SEC_TYPE="ext2"
```

4. Identify the UUID number for the LUN:

```
dumpe2fs device_path_name | grep UUID
```

5. Identify the directory on which the DMMP device is mounted:

```
df -h
```

In the following example, /mnt/dmmp_ext3 represents the directory on which the DMMP device is mounted:

```
[root@IBMX3550M3-229-169 ~]# df -h
Filesystem Size Used Avail Use% Mounted on
/dev/mapper/dmmp_fs_lun
1008M 34M 924M 4% /mnt/dmnp_ext3
```

6. Verify in the /etc/fstab file that the mount points for the DMMP device are defined:

```
cat /etc/fstab
```

The DMMP device name and mount directory should be displayed in the output.

7. Create a backup of the /etc/fstab file:

```
cp /etc/fstab /etc/fstab_pre_transition_bkup
```

8. Edit the /etc/fstab file to replace the DMMP device name with its respective file system UUID number.

Testing DMMP devices with file systems 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 of your Red Hat Enterprise Linux (RHEL) host, you can test your transitioned clustered Data ONTAP LUNs to verify that you can mount your DMMP device 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 your test host and your LUNs must be ready for transition.

You should maintain hardware parity between the test host and the source host.

Perform these steps on the test host.

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. Obtain the new SCSI device names for the clustered Data ONTAP LUNs:

```
sanlun lun show
```

In the following example, /dev/sdl is the SCSI device name for the lun_dmmp_raw LUN, and /dev/sdk is the SCSI device name for the lun_dmmp_raw_alias LUN:

```
[root@ibmx3550-229-108 /]# sanlun lun show
controller(7mode/E-Series)/
vserver (cDOT/FlashRay) lun-pathname          filename
-----
vs_brb   /vol/dmmp_raw_vol/lun_dmmp_raw             /dev/sdl
vs_brb   /vol/dmmp_raw_alias_vol/lun_dmmp_raw_alias  /dev/sdk
```

4. Configure the DMMP devices for your clustered Data ONTAP LUNs:

```
multipath
```

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

```
multipath -ll
```

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

6. Identify the file system configured on the DMMP device:

```
blkid | grep -i device_handle_ID
```

7. Determine whether a mount point entry for the logical volume exists in the /etc/fstab file on the source host.
8. If a mount point entry exists for the logical volume on the source host, manually edit the /etc/fstab file on the test host to add the mount point entries.
9. Mount the LUN:

```
mount -a
```

10. Verify that the DMMP device is mounted:

```
mount
```

11. Perform your testing as needed.

12. After you have completed your testing, shut down the test host:

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

13. 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 the cutover phase when transitioning LUNs with mount points using DMMP device names on Linux hosts](#)

Preparing for the cutover phase when transitioning LUNs with mount points using DMMP device names on Linux hosts

If you are transitioning a LUN with a mount point using an alias name on a Linux host, there are several steps you must perform before entering 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.

- 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 Systems operation in the 7MTT.

Steps

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

```
umount dir_name
```

4. Flush the 7-Mode LUN DMMP device ID:

```
multipath -f device_name
```

If needed, you can get the DMMP device name from the **OS Device ID** column under the **SAN Host LUNs** tab in the *Inventory Assessment Workbook*.

Related information

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

Remounting DMMP devices on Linux hosts after transition

After transition from ONTAP operating in 7-Mode to clustered Data ONTAP, you must remount your DMMP devices for RHEL 5 and RHEL 6. Your 7-Mode LUNs are not accessible to the host until your DMMP devices are mounted.

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, 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 is complete.

Steps

1. Generate the 7-Mode to 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 project-name -s svm-name -o output-file
```

For example:

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



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

2. Make a note of the new ONTAP LUN device handle ID from the LUN mapping file.

3. 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.

4. Discover new ONTAP LUNs:

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

5. Verify that the ONTAP LUNs are discovered:

```
sanlun lun show
```

The ONTAP LUN's SCSI devices should be listed under the filename column.

6. Configure DMMP devices for ONTAP LUNs:

```
multipath
```

7. Verify that the DMMP devices are present:

```
multipath -ll LUN_SCSI_device_name
```

In the following example, 3600a098051764937303f4479515a7451 represents the DMMP device handle ID:

```
[root@IBMX3550M3-229-169 ~]#multipath -ll /dev/sdq
3600a098051764937303f4479515a7451 dm-6 NETAPP,LUN C-Mode
```

8. Mount the LUN:

```
*mount device_name mountpoint
```

If the mount points are defined in the `/etc/fstab` file, you can run the `mount -a` command to mount all of the mount points.

9. Verify the mount points:

```
mount
```

Transitioning LUNs with mount points using DMMP alias names

When you transition a LUN with a mount point using an alias name, you must perform specific steps to prepare for the cutover phase and you must remount the LUNs after the transition.

Related information

[Preparing for the cutover phase when transitioning LUNs with mount points using DMMP device names on Linux hosts](#)

[Remounting LUNs with mount points using DMMP alias names on Linux hosts after transition](#)

Preparing RHEL LUNs with mount points using DMMP alias names for transition using the CLI

Before transition of a mount point using a DMMP device name, you must replace the DMMP device name with its respective file system UUID number.

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

Steps

1. Identify and record the SCSI device ID for the LUN to be transitioned:

```
sanlun lun show
```

The SCSI device ID is listed under the filename column in the output.

2. Identify and record the DMMP device name for the LUN to be transitioned:

```
multipath -ll SCSI_device_ID
```

In the following example, 360a9800037534562572b453855496b41 is the DMMP device name:

```
[root@IBMx3550M3-229-169 ~]# multipath -ll /dev/sdc
dmmp_fs_lun (360a9800037534562572b453855496b41) dm-3 NETAPP, LUN
[size=1.0G] [features=3 queue_if_no_path pg_init_retries 50]
[hwhandler=0][rw]
\_ round-robin 0 [prio=2][active]
  \_ 9:0:0:1 sdc 8:32 [active][ready]
  \_ 9:0:0:1 sdg 8:96 [active][ready]
```

3. Identify the file system configured on the DMMP device:

```
blkid | grep -i DMMP_device_name
```

The TYPE value in the output identifies the file system.

In the following example, the file system is ext3.

```
[root@ibmx3550-229-108 ~]# blkid | grep -i
3600a09804d532d79565d47617679658
/dev/mapper/3600a09804d532d79565d47617679658:
UUID="450b999a-4f51-4828-8139-29b20d2f8708" TYPE="ext3" SEC_TYPE="ext2"
```

4. Identify the UUID number for the LUN:

```
dumpe2fs device_path_name | grep UUID
```

5. Identify the directory on which the DMMP device is mounted:

```
df -h
```

In the following example, /mnt/dmmp_ext3 represents the directory on which the DMMP device is mounted:

```
[root@IBMX3550M3-229-169 ~]# df -h
Filesystem Size Used Avail Use% Mounted on
/dev/mapper/dmmp_fs_lun
1008M 34M 924M 4% /mnt/dmnp_ext3
```

6. Verify in the `/etc/fstab` file that the mount points for the DMMP device are defined:

```
cat /etc/fstab
```

The DMMP device name and mount directory should be displayed in the output.

7. Create a backup of the `/etc/fstab` file:

```
cp /etc/fstab /etc/fstab_pre_transition_bkup
```

8. Edit the `/etc/fstab` file to replace the DMMP device name with its respective file system UUID number.

Testing LUNs with mount points using DMMP alias names 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 of your Red Hat Enterprise Linux (RHEL) host, you can test your transitioned clustered Data ONTAP LUNs with mount points using alias names 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 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 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.

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. Obtain the new SCSI device names for the clustered Data ONTAP LUNs:

```
sanlun lun show
```

In the following example, `/dev/sdl` is the SCSI device name for the `lun_dmmp_raw` LUN, and `/dev/sdk` is the SCSI device name for the

`lun_dmmp_raw_alias` LUN:

```
[root@ibmx3550-229-108 /]# sanlun lun show
controller(7mode/E-Series)/
vserver (cDOT/FlashRay)      lun-pathname      filename
-----
vs_brb      /vol/dmmp_raw_vol/lun_dmmp_raw      /dev/sdl
vs_brb      /vol/dmmp_raw_alias_vol/lun_dmmp_raw_alias  /dev/sdk
```

4. Configure the DMMP devices for your clustered Data ONTAP LUNs:

multipath

5. Obtain the device handle IDs for the clustered Data ONTAP LUNs:

multipath -ll

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

6. Check whether an alias is defined in the `/etc/multipath.conf` file on the source host.
7. Manually copy the alias configuration to the `/etc/multipath.conf` file on the test host, but replace the 7-Mode device handle ID with the corresponding clustered Data ONTAP device handle ID.
8. Use the `multipath` command to configure DMMP devices for your clustered Data ONTAP LUNs.
9. Identify the file system created on the DMMP alias device:

blkid dmmp_device_name

10. Mount the DMMP device:

mount

11. Perform your testing as needed.
12. After you have completed your testing, shut down the test host:

shutdown -h -t0 now

13. 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 the cutover phase when transitioning LUNs with mount points using DMMP device names on Linux hosts](#)

Preparing for the cutover phase when transitioning LUNs with mount points using DMMP device names on Linux hosts

If you are transitioning a LUN with a mount point using an alias name on a Linux host, there are several steps you must perform before entering 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.

- 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 Systems operation in the 7MTT.

Steps

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

```
umount dir_name
```

4. Flush the 7-Mode LUN DMMP device ID:

```
multipath -f device_name
```

If needed, you can get the DMMP device name from the **OS Device ID** column under the **SAN Host LUNs** tab in the *Inventory Assessment Workbook*.

Related information

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

Remounting LUNs with mount points using DMMP alias names on Linux hosts after transition

After transition from ONTAP operating in 7-Mode to clustered Data ONTAP, you must remount your LUNs with mount points. The 7-Mode volumes are offline and the 7-Mode LUNs are not 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 7MTT.
- For CFTs, perform these steps after the Import Data & Configuration operation in the 7MTT.
 1. Generate the 7-Mode to 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 *project-name* -s *svm-name* -o *output-file*

For example:

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



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

2. Make a note of the ONTAP device handle ID in the LUN mapping file.

3. 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.

4. Discover the new ONTAP LUNs:

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

5. Verify that the ONTAP LUNs are discovered:

```
sanlun lun show
```

The ONTAP LUN's SCSI devices should be listed in the `device filename` column.

An example of a SCSI device name is `/dev/sdp`.

6. In the `/etc/multipath.conf` file, replace the 7-Mode device handle ID with the clustered Data ONTAP LUN's device handle ID so that the `alias` name points to the clustered Data ONTAP LUN ID.

You should update the `multipaths` section as displayed below. The following example shows the `/etc/multipath.conf` file, before replacing the 7-Mode LUN ID. In this example, the LUN ID `360a9800037534562572b453855496b43` is pointing to the `dmmp_fs_lun` alias name.

```

multipaths {
    multipath {
        wwid      360a9800037534562572b453855496b43
        alias      dmmp_fs_lun
    }
}

```

After replacing the 7-Mode LUN ID with the ONTAP LUN ID
360a9800037534562572b453855496b43, the example file looks as follows:

```

multipaths {
    multipath {
        wwid      3600a098051764937303f4479515a7452
        alias      dmmp_fs_lun
    }
}

```

7. Configure DMMP devices for ONTAP LUNs:

`multipath`

8. Verify that the DMMP alias points to the ONTAP LUN device handle ID:

`multipath -ll device_handle_ID`

9. Mount the ONTAP LUN to its mount point directory:

`mount /dev/mapper/alias_name mount_dir_name`

If the mount points are defined in the `/etc/fstab` file, use the `mount -a` command to mount the LUN.

10. Verify that the DMMP device is mounted:

mount

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_project-name_ -s svm-name -o output-file
```

For example:

```
transition cft export lunmap -p SanWorkLoad -s svml -o 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_name mount_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
```

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/LogVol101 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/LogVol101 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 © 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.