




7-Mode Transition Tool 2.1

SAN Host Transition and Remediation Supplemental Guide



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-10061_A0
June 2015

Table of Contents

About the <i>SAN Host Transition and Remediation Supplemental Guide</i>	4
Introduction to VMware ESXi host remediation.....	4
1.1 What is supported in ESXi with 7MTT SAN Transition	4
1.2 Pretransition stage.....	4
1.2.1 Pretransition prerequisites.....	4
1.3 Pretransition steps for VMware vSphere ESXi versions	5
1.4 Pretransition check for VM Snapshot.....	10
1.5 Transition stage	12
1.6 ESXi Posttransition Stage.....	12
1.6.1 Remediation of SAN Boot hosts.....	13
1.6.2 VMFS Datastore Remediation.....	13
1.6.3 Remediation of Raw Devices Mapped LUNs (RDM) attached to VMs	18
1.6.4 Enabling Compare and Write (CAW) on ESX	27
1.7 Configure recommended settings for ESXi host	30
RHEL 5 series.....	31
Data LUN Scenarios.....	31
1.8 DMMP device without file system	31
1.8.1 Pretransition	31
1.8.2 Before entering cutover phase	32
1.8.3 Posttransition	33
1.9 Mount point using DMMP device name	36
1.9.1 Pretransition	36
1.9.2 Before entering cutover phase	38
1.9.3 Posttransition	38
1.10 Mount point using alias name	40
1.10.1 Pretransition.....	40
1.10.2 Before entering cutover phase.....	42
1.10.3 Posttransition	42
1.11 File system on LVM device	44

1.11.1	Pretransition.....	44
1.11.2	Before entering cutover phase.....	47
1.11.3	Posttransition	48
SAN Boot LUN Scenarios		53
1.12	Protocol: FC/FCoE.....	53
1.12.1	Pretransition.....	53
1.12.2	Before entering cutover phase.....	56
1.12.3	Posttransition	56
RHEL 6 series.....		58
Data LUN Scenarios		58
1.13	DMMP device without file system	58
1.13.1	Pretransition.....	58
1.13.2	Before entering cutover phase.....	59
1.13.3	Posttransition	60
1.14	DMMP device with file system	62
1.14.1	Before entering cutover phase.....	65
1.14.2	Posttransition	66
1.15	File system on LVM device	67
1.15.1	Pretransition.....	67
1.15.2	Before entering cutover phase.....	69
1.15.3	Posttransition	70
SAN Boot LUN Scenarios		74
1.16	Protocol: FC/FCoE.....	74
1.16.1	Pretransition.....	74
1.16.2	Before entering cutover phase.....	77
1.16.3	Posttransition	77
1.17	Protocol: iSCSI	78
1.17.1	Pretransition.....	78
1.17.2	Before entering cutover phase.....	83
1.17.3	Posttransition	84
Windows host transition steps		85
1.18	Pretransition.....	85

1.19 Transition.....	86
1.20 Posttransition.....	88
1.21 Exceptions and known issues.....	90
Copyright information.....	91
Trademark information.....	92
How to send comments about documentation and receive update notification.....	93
Copyright information.....	88
Trademark information.....	89
How to send comments about documentation and receive update notification.....	90

About the *SAN Host Transition and Remediation Supplemental Guide*

The *SAN Host Transition and Remediation Supplemental Guide* provides the pretransition and posttransition procedures for VMware ESXi, Red Hat Linux, and Windows hosts when transitioning data and configurations to clustered Data ONTAP 8.3 and later using the 7-Mode Transition Tool.

Introduction to VMware ESXi host remediation

This section describes VMware ESXi host remediation required for SAN transition as part of the transition to clustered Data ONTAP 8.3 and later using **7-Mode Transition Tool (7MTT)**. The document details ESXi version support, feature support, guest OS dependencies, and actions required on the host to avoid transition failures.

1.1 What is supported in ESXi with 7MTT SAN Transition

The following ESXi versions and features are supported for SAN transition using the 7-Mode Transition Tool:

- ESXi 5.0, 5.1 and 5.5 all updates
- VMFS3 and VMFS5 datastores
- SAN boot configuration
- RDM (RAW device map) devices
- All guest OS based on the support matrix in [IMT](#)
- All SAN protocols (FC/FCoE/iSCSI) are supported based on the support matrix in [IMT](#)

The transition to clustered Data ONTAP SAN using 7MTT requires a cutover stage where, if the host is SAN booted or if SAN booted LUN are part of the transition, the ESXi host experiences a down time. As an alternative, customers who cannot take down time can use **VMware Storage vMotion** to migrate VMs to clustered Data ONTAP LUNs, provided they do not have RAW device mapping (RDM) on any Data ONTAP LUNs in their VMware vSphere infrastructure. For more information, see [Migrate a Virtual Machine with Storage vMotion in the vSphere Client on the VMware website](#).

1.2 Pretransition stage

The following information is a list of pretransition prerequisites that must be completed before initiating transition using the 7MTT tool.

1.2.1 Pretransition prerequisites

The following requirements must be verified and completed before transitioning hosts using the 7-Mode Transition Tool:

1. Clustered Data ONTAP 8.3 and later should be configured and set up as described in the *7-Mode Transition Tool Guide*, available on the [NetApp Support Site](#).
2. Manually compile a list of ESXi host's IP address, hostname, and authentication details of hosts that are part of the 7-Mode Transition Tool SAN transition. You use this information to perform remediation on ESXi hosts.

3. Zoning between FC/FCoE hosts and new clustered Data ONTAP controllers must be completed as described in the 7-Mode Transition Tool Guide, available on the [NetApp Support Site](#).
4. Collect configuration details of ESXi hosts and 7-Mode controllers using the *Collect & Assess* option available in the 7-Mode Transition Tool.
5. The version of Data ONTAP operating in 7-Mode and connected ESXi hosts, along with the HBA driver and firmware, must be a supported configuration for clustered Data ONTAP. Check the [Interoperability Matrix Tool \(IMT\)](#) for the ESXi and clustered Data ONTAP support matrix.

If ESXi hosts are not a supported configuration for clustered Data ONTAP, you must reconfigure the host before initiating transition of data to clustered Data ONTAP.

6. If high availability (HA) and Distributed Resource Scheduler (DRS) are configured, they must be disabled.

1.3 Pretransition steps for VMware vSphere ESXi versions

ESXi 4.x and earlier

Data ONTAP 8.3 and later supports ESXi 5.0. Earlier versions of ESX and ESXi are not supported; therefore, the transition plan does not support ESXi/ESX 4.x and earlier. To migrate configurations with ESXi4.x from Data ONTAP operating in 7-Mode to clustered Data ONTAP 8.3 and later, you must first upgrade those hosts to a supported clustered Data ONTAP 8.3 and later configuration. For details see the following sections.

ESXi 5.x

ESXi 5.0, 5.1 and 5.5 and all updates are included in the transition plan according to IMT support for clustered Data ONTAP 8.3 and later and the source Data ONTAP operating in 7-Mode version from which you are transitioning.

Examples:

If you are planning to transition from Data ONTAP 8.1.4 operating in 7-Mode to clustered Data ONTAP 8.3 and later, the ESXi host connected to the source running Data ONTAP 8.1.4 operating in 7-Mode must have an ESXi version supported by the destination running clustered Data ONTAP 8.3 and later system, otherwise, the ESXi host must be upgraded to a version supported by clustered Data ONTAP 8.3 and later before the transition. Refer to the [IMT](#) for support details.

If the host is an ESXi4.1U3 and transition is from Data ONTAP 8.1.4 operating in 7-Mode to clustered Data ONTAP 8.3 and later, ESXi4.1U3 is not supported by clustered Data ONTAP 8.3 and later. By referring to the [IMT](#), you find that ESXi5.0U3 is supported by both Data ONTAP 8.1.4 operating in 7-Mode and clustered Data ONTAP 8.3 and later. Therefore, you can upgrade the host to ESXi5.0U3 (Refer to the [Upgrade VMware KB](#)) before transition so that it is compatible for SAN transition from Data ONTAP 8.1.4 operating in 7-Mode to clustered Data ONTAP 8.3 and later.

Original source configuration	Upgrade ESXi	Migrate to clustered Data ONTAP
Data ONTAP 8.1.4 operating in 7-Mode	Data ONTAP 8.1.4 operating in 7-Mode	Clustered Data ONTAP 8.3 and later
ESXi4.1U3	ESXi5.0U3	ESXi5.0U3

If the host is an ESXi4.0U4 and transition is from Data ONTAP 7.3.7 operating in 7-Mode to clustered Data ONTAP 8.3 and later, ESXi4.0U4 is not supported by clustered Data ONTAP 8.3 and later and Data ONTAP 7.3.7 operating in 7-Mode is not compatible for SAN transition using 7MTT. In this case, you must upgrade Data ONTAP to a 7MTT SAN transition compatible version and then upgrade ESXi to a clustered Data ONTAP 8.3 and later supported version. Data ONTAP 8.1.4 operating in 7-Mode is compatible with 7MTT SAN Transition, and by referring to the [IMT](#), you determine that ESXi5.0U3 is supported by both Data ONTAP 8.1.4 operating in 7-Mode and clustered Data ONTAP 8.3 and later. Upgrading to Data ONTAP 8.1.4 operating in 7-Mode and then upgrading the host to ESXi5.0U3 before transition makes the configuration compatible for 7MTT SAN transition. Refer to the [Upgrade VMware KB](#).

Original source configuration	Upgrade Data ONTAP	Upgrade ESXi	Migrate to clustered Data ONTAP
Data ONTAP 7.3.7 operating in 7-Mode	Data ONTAP 8.1.4 operating in 7-Mode	Data ONTAP 8.1.4 operating in 7-Mode	Clustered Data ONTAP 8.3 and later
ESXi4.0U4	ESXi4.0U4	ESXi5.0U3	ESXi5.0U3

Pretransition HBA, Driver, and Firmware Support Check

NetApp supports only transition of software iSCSI initiator in ESXi 5.x.

For FC and FCoE, only QLogic and Emulex initiators in ESXi are supported by Netapp. ESXi QLogic and Emulex initiator driver and firmware support is available as described in the [IMT](#).

For example, if you are transitioning from Data ONTAP 8.1.4 operating in 7-Mode to clustered Data ONTAP 8.3 and later, the ESXi host connected to the Data ONTAP 8.1.4 source must have an FC and FCoE initiator driver and firmware version supported by the clustered Data ONTAP 8.3 and later destination. If the ESXi FC or FCoE initiator driver and firmware is not supported, you must upgrade to a clustered Data ONTAP 8.3 and later supported version before transition. See the support matrix in the [IMT](#).

If the host is ESXi5.0U3 with QLogic SANblade QLE2562 (Driver: 934.5.29.0-1vmw (Async), Firmware: 7.00.02) and transition is from Data ONTAP 8.1.4 operating in 7-Mode to clustered Data ONTAP 8.3 and later, the ESXi5.0U3 QLogic SANblade QLE2562 Firmware: 7.00.02 is not supported by clustered Data ONTAP 8.3 and later. But if you refer to [IMT](#), you can find that ESXi5.0U3 QLogic SANblade QLE2562 Firmware: 7.01.00 is supported by clustered Data ONTAP 8.3 and later, and you can upgrade the QLogic SANblade QLE2562 firmware on the ESXi5.0U3 host before transition, making it compatible for SAN transition from Data ONTAP 8.1.4 operating in 7-Mode to clustered Data ONTAP 8.3 and later.

Original source configuration	Upgrade firmware	Migrate to clustered Data ONTAP
Data ONTAP 8.1.4 operating in 7-Mode	Data ONTAP 8.1.4 operating in 7-Mode	Clustered Data ONTAP 8.3 and later
ESXi5.0U3 QLogic SANblade QLE2562 (Driver: 934.5.29.0-1vmw (Async), Firmware: 7.00.02)	ESXi5.0U3 QLogic SANblade QLE2562 (Driver: 934.5.29.0-1vmw (Async), Firmware: 7.01.00)	ESXi5.0U3 QLogic SANblade QLE2562 (Driver: 934.5.29.0-1vmw (Async), Firmware: 7.01.00)

Pretransition check for VM Guest OS

1. Verify that ESXi hosts and Data ONTAP operating in 7-Mode information is collected via **Collect and Assess** option of 7MTT tool.

2. Guest OS in the VMs would require additional pre-remediation steps before transition of LUN's from Data ONTAP operating in 7-Mode to clustered Data ONTAP 8.3 and later. Refer to below steps to complete the pre-requisite remediation procedure for specific guest OS

LINUX Guest OS

If 7-Mode LUN's are mapped as Physical compatible RDM (PTRDM) to Linux VMs as a boot device then it would require pre remediation changes.

1. Check **/boot/grub/menu.lst** to determine whether **kernel** or **initrd** parameters contain any reference to SCSI device/partition serial number.

For example:

```
# cat /boot/grub/menu.lst

# Modified by YaST2. Last modification on Fri Oct 17 02:08:40 EDT 2014
default 0
timeout 8
##YaST - generic_mbr
gfxmenu (hd0,1)/boot/message
##YaST - activate
###Don't change this comment - YaST2 identifier: Original name: linux###
title SUSE Linux Enterprise Server 11 SP2 - 3.0.13-0.27
    root (hd0,1)
    kernel /boot/vmlinuz-3.0.13-0.27-default root=/dev/disk/by-id/scsi-360a9800032466879362b45777447462d-part2 resume=/dev/disk/by-id/scsi-360a9800032466879362b45777447462d-part1 splash=silent crashkernel=128M-:64M showopts
    vga=0x314
    initrd /boot/initrd-3.0.13-0.27-default
```

In the previous example root=**/dev/disk/by-id/scsi-360a9800032466879362b45777447462d-part2** resume=**/dev/disk/by-id/scsi-360a9800032466879362b45777447462d-part1** refers to the SCSI device serial number as device path.

2. You must update the grub boot menu file to look in the UUID-based device path instead of serial number-based device path. Use the following command to get the UUID of the device/partition:

```
#blkid
```

For example:

The first relationship/mapping between the SCSI serial number, SCSI device/partition or serial number can be determined using the following command:

```
# ls -l /dev/disk/by-id
total 0
lrwxrwxrwx 1 root root  9 Oct 27 06:54 edd-int13_dev80 -> ../../sda
lrwxrwxrwx 1 root root 10 Oct 27 05:09 edd-int13_dev80-part1 -> ../../sda1
lrwxrwxrwx 1 root root  9 Oct 27 06:54 scsi-360a9800032466879362b45777447462d -> ../../sda
lrwxrwxrwx 1 root root 10 Oct 27 05:09 scsi-360a9800032466879362b45777447462d-part1 -
> ../../sda1
lrwxrwxrwx 1 root root 10 Oct 27 02:21 scsi-360a9800032466879362b45777447462d-part2 -
> ../../sda2
lrwxrwxrwx 1 root root  9 Oct 27 06:54 scsi-360a9800032466879362b457774474638 -> ../../sdb
lrwxrwxrwx 1 root root  9 Oct 27 06:54 wwn-0x60a9800032466879362b45777447462d -> ../../sda
lrwxrwxrwx 1 root root 10 Oct 27 05:09 wwn-0x60a9800032466879362b45777447462d-part1 -
> ../../sda1
lrwxrwxrwx 1 root root 10 Oct 27 02:21 wwn-0x60a9800032466879362b45777447462d-part2 -
> ../../sda2
lrwxrwxrwx 1 root root  9 Oct 27 06:54 wwn-0x60a9800032466879362b457774474638 -> ../../sdb
```

```
#ls -l /dev/disk/by-uuid/
total 0
lrwxrwxrwx 1 root root 10 Oct 27 02:21 33d43a8b-cfae-4ac4-9355-36b479cfa524 -> ../../sda2
lrwxrwxrwx 1 root root 9 Oct 27 06:54 603e01f8-7873-440a-9182-878abff17143 -> ../../sdb
lrwxrwxrwx 1 root root 10 Oct 27 05:09 c50b757b-0817-4c19-8291-0d14938f7f0f -> ../../sda1
```

From the blkid output, match the UUID with the SCSI device or serial number:

```
#blkid
/dev/sda1: UUID="c50b757b-0817-4c19-8291-0d14938f7f0f" TYPE="swap"
/dev/sda2: UUID="33d43a8b-cfae-4ac4-9355-36b479cfa524" TYPE="ext3"
/dev/sdb: UUID="603e01f8-7873-440a-9182-878abff17143" SEC_TYPE="ext2" TYPE="ext3"
```

3. Use the UUID you just retrieved to update the device reference in the grub boot menu.lst file.

For example:

After the update above /boot/grub/menu.lst would look like this:

```
# Modified by YaST2. Last modification on Fri Oct 17 02:08:40 EDT 2014
default 0
timeout 8
##YaST - generic_mbr
gfxmenu (hd0,1)/boot/message
##YaST - activate
###Don't change this comment - YaST2 identifier: Original name: linux###
title SUSE Linux Enterprise Server 11 SP2 - 3.0.13-0.27
    root (hd0,1)
    kernel /boot/vmlinuz-3.0.13-0.27-default root=/dev/disk/by-uuid/e5127cdf-8b30-
418e-b0b2-35727161ef41 resume=/dev/disk/by-uuid/d9133964-d2d1-4e29-b064-7316c5ca5566
    splash=silent crashkernel=128M-:64M showopts vga=0x314
    initrd /boot/initrd-3.0.13-0.27-default
```

4. Similarly, update **/etc/fstab** for any references to the device path based on serial number. If the swap partition is referred by device serial number, you must create a label as shown in the example below:

```
/dev/disk/by-id/scsi-360a9800032466879362b45777447462d-part1 swap swap
defaults 0 0
/dev/disk/by-id/scsi-360a9800032466879362b45777447462d-part2 / ext3
acl,user_xattr 1 1
proc /proc proc defaults 0 0
sysfs /sys sysfs noauto 0 0
debugfs /sys/kernel/debug debugfs noauto 0 0
devpts /dev/pts devpts mode=0620,gid=5 0 0
```


In **/etc/fstab** replace the reference to scsi serial number with UUID noted from blkid command.

The modified **/etc/fstab** file would look like the following:

```
cat /etc/fstab
UUID="c50b757b-0817-4c19-8291-0d14938f7f0f swap swap defaults
0 0
UUID="33d43a8b-cfae-4ac4-9355-36b479cfa524 / ext3 acl,user_xattr
1 1
proc /proc proc defaults 0 0
sysfs /sys sysfs noauto 0 0
debugfs /sys/kernel/debug debugfs noauto 0 0
devpts /dev/pts devpts mode=0620,gid=5 0 0
```

WINDOWS Guest OS

If Windows VMs use Physical compatible RDM (PTRDM) devices, before transition, take the disks offline on the Windows VM using Disk Manager.




1.4 Pretransition check for VM Snapshot

VM snapshots of VMs with only VMFS vDisks and Physical RDM (PT RDM) will work out of the box after 7MTT SAN transition. But VM Snapshots of VMs with any Virtual RDM (NPTRDM) will break after transition. Therefore, as a pretransition step, all snapshots of VMs with Virtual RDM attached must be removed. Identify all VMs with Virtual RDMs attached using Inventory Assessment Workbook collected by 7MTT.


In the Inventory Assessment Workbook generated by ICT, navigate to the "Host VMs" tab and verify the contents of the "VM Snapshots" and "NPTRDM" columns. If there are any snapshots listed under the "VM Snapshots" column and the "NPTRDM" column that have a value greater than 0, this indicates the VM has a Virtual RDM (NPTRDM) and that it includes VM snapshots. Note the VM name and VMID and remove all snapshots for these VMs using the following steps.

Removing VM Snapshots Using vSphere Client:

1. Using vSphere Client, open the ESXi host or vCenter server managing the ESXi host.
2. Right-click the VM which needs the snapshot removed, and from the menu select *Snapshot->Snapshot Manager* to open Snapshots window.



3. If snapshots for the VM are listed in the Snapshots window, click *Delete All* to delete all snapshots.



Removing VM snapshots using ESXi CLI

1. Login to ESXi console using SSH.

- Using the **VMID** that was noted previously from the ICT Inventory Assessment Workbook “Host VMs” tab, delete the VM snapshots.
- Use the following command to remove all VM snapshots of the VM with **VMID**:

```
# vim-cmd vmsvc/snapshot.removeall <VMID>
```

Example: If VMID is 22 then


```
# vim-cmd vmsvc/snapshot.removeall 22
```

1.5 Transition stage

- After ESXi pretransition checks are complete, the LUN transition from Data ONTAP operating in 7-Mode to clustered Data ONTAP 8.3 and later can be initiated using the 7-Mode Transition Tool. ESXi hosts, VMs, and applications can be online until pre-cutover phase of 7MTT transition.
- Plan your downtime. Ensure that all applications and Guest operating systems are powered down at pre-cutover phase. If ESXi hosts are SAN booted, or if SAN booted LUNs are part of the transition, ESXi hosts must also be powered down at pre-cutover phase of the 7MTT transition.

If the host is not SAN booted and you do not plan to shut it down, you might experience the all-paths-down (ADP) condition. To avoid the ADP condition, you can mask 7-Mode LUNs using the steps described in [VMware knowledge base article 1009449](#).

- After the VMs and ESXi hosts are powered off, complete the 7MTT transition as illustrated in the 7MTT transition dashboard.



1.6 ESXi posttransition Stage

After the 7MTT transition is complete, perform the following posttransition checks:

- Generate a LUN mapping file. This file is important for ESXi posttransition remediation because it provides the mappings between 7-Mode array to clustered Data ONTAP array LUNs. You can generate the file using the following 7MTT CLI command:

```
C:\Program Files\NetApp\operating in 7-Mode Transition Tool\bin> transition export lunmap -s <transition session-name> -o <filename>.csv
```

- Verify igroup and initiator mappings. 7MTT recreates the same igroup with initiators that were used in Data ONTAP operating in 7-Mode, and remaps the clustered Data ONTAP LUN to the host.
- Ensure zoning is appropriate for the new clustered Data ONTAP 8.3 and later target before rediscovering clustered Data ONTAP 8.3 and later LUNs on the ESXi host. For more information, see the *7-Mode Transition Tool Guide*, available on the [NetApp Support Site](#).

1.6.1 Remediation of SAN Boot hosts

- If the ESXi host was configured for SAN boot, then reconfigure FC/FCoE HBA BIOS to boot from the SAN boot LUN of the Clustered Data ONTAP system.
- Once the ESXi host is up and running as expected, you must remediate the host based on the configurations you had before the transition.
- If an ESXi 5.x host was accessing the 7-Mode system with the VMware Software iSCSI initiator, then the host requires a reconfiguration of the Software iSCSI initiator and dependent ports.

To perform the reconfiguration task, enable the reconfigured host Software iSCSI initiator to discover the new clustered Data ONTAP 8.3 and later target, retrieve the target iSCSI IP and IQN used by the new clustered Data ONTAP 8.3 and later target. See the *Clustered Data ONTAP 8.3 Administration Guide* on the [NetApp Support Site](#).

To make the required changes on the Software iSCSI initiator on VMware ESXi hosts, see the VMware [vSphere ESXi5.x Storage Guide](#) (Chapter 11 - Configuring iSCSI Adapters and Storage).

Note: If the target IP subnet has changed, then the respective IP subnet changes are required on the host iSCSI initiator ports as well.

- VMFS Datastores created from the boot LUN in a default installation must be remounted. Follow the steps in [VMFS Datastore Remediation](#) section to perform this task.

1.6.2 VMFS Datastore Remediation


After transition to clustered Data ONTAP, **VMFS3** and **VMFS5** will be discovered as VMFS volumes on a snapshot LUN because of a LUN serial number change that occurred during transition. This is a result of the VMFS signature in the VMFS filesystem being intact on the new clustered Data ONTAP LUN. Remounting the volumes persistently without altering or resignaturing the VMFS signature brings the datastores back intact, which also brings all the VMs back to their original state. The exceptions are VMs with RDM LUN which have lost their RDM mapping. See the RDM section for details.

- In the Inventory Assessment Workbook that is generated by ICT, navigate to the "SAN Host Filesystems" tab and then verify in the column "Drive / Mount / Datastore Name" the filesystems and datastores that were mounted on the host before transition. Also note the corresponding underlying lun naa ID's from the "SCSI Device ID/Device Name" column for that datastore.
- Verify that all the naa ID's noted for a given datastore are listed in the 7MTT mapping file that was generated after the transition. Also, verify that all the noted naa ID's have a new, corresponding naa ID in the "LUN WWID" column of the 7MTT mapping file.
 - If none of the naa ID's are present in the 7MTT mapping file, this indicates that the datastore and its underlying LUNs were not part of the 7MTT transition and no remediation is required for this datastore. Proceed to next datastore in the Inventory Assessment Workbook.


- o If only partial naa ID's of a datastore are found in the 7MTT mapping file then you cannot remediate the datastore. Your transition is incomplete because only a partial set of LUNs underlying this datastore are transitioned. You cannot proceed any further with remediation.
- o If all naa ID's of a datastore are found in the 7MTT mapping file, you can remediate the datastore using the following steps:

Remounting VMFS Volumes after transition using vSphere Client:


1. Login to the ESXi host or the vCenter managing the ESXi host using vSphere client.
2. Under Hosts and Clusters, select the ESXi host.
3. Navigate to *Configuration->Hardware->Storage* to add the datastores to your datastore list.
4. Click *Add storage* to open the *Add Storage* window. Select the *Disk/LUN* option, and then click *Next*.



5. In the displayed list of LUNs, locate the **VMFS_label** column displaying the name of the datastore.



6. Select the LUN you identified earlier and complete the wizard, selecting the option **Keep the existing signature** in the Select VMFS Mount Options window.



7. Repeat these steps for all the LUNs displaying the datastore name in the **VMFS_label** column.
8. All datastores will be remounted and all VMs will be active.

Remounting VMFS Volumes after transition using ESXi CLI/Console:

1. Login to ESXi console using SSH.
2. From either the ESXi console or CLI, list all the newly added LUNs with existing vmfs signature and VMFS label, using the following command:

```
# esxcfg-volume -l

VMFS UUID/label: 53578567-5b5c363e-21bb-001ec9d631cb/datastore1
Can mount: Yes
Can resignature: Yes
Extent name: naa.600a098054314c6c445d446f79716475:1      range: 0 - 409599
(MB)
```

3. Use the `esxcfg-volume -M|--persistent-mount <VMFS UUID|label>` command to remount the vmfs volumes persistently with same signature provided the original 7-Mode LUN is unmapped from the host or is not online.

```
# esxcfg-volume -M datastore1
```


For more details about remounting, see the VMware KB: [vSphere handling of LUNs detected as snapshot LUNs](#)

Remounting spanned VMFS Volumes after transition using vSphere Client:

An administrator must ensure that a VMFS volume spanning multiple LUNs has all of its spanned LUNs transitioned to the destination Data ONTAP system.

To remount Spanned VMFS volume with existing signature, you can follow the VMFS remounting steps described earlier. The LUNs on the VMFS span volume are listed as shown in the figure shown earlier of a **vSphere Client** during remount operation.

All LUNs that belong to a Spanned VMFS volume will have the same VMFS label as shown in the earlier “Select Disk/LUN” figure, and the first LUN in the span, circled in red, will be marked **(head)**. You must select the head LUN to complete the remount operation.



Remounting spanned VMFS Volumes after transition using ESXi CLI/Console:

1. Login to ESXi console using SSH.
2. When listed from ESXi CLI Spanned VMFS volumes will display as follows:

```
# esxconfig-volume -l

VMFS UUID/label: 537dedf6-1f1a994b-ab4d-001ec9d631cb/vmfs-span-ds
Can mount: No (the original volume has some extents online)
Can resignature: Yes
Extent name: naa.600a098054314c6c445d446f79716431:1      range: 0 - 9983
(MB)
Extent name: naa.600a098054314c6c445d446f79716432:1      range: 9984 -
19967 (MB)
Extent name: naa.600a098054314c6c445d446f79716433:1      range: 19968 -
29951 (MB)
```

3. You can use the `esxcfg-volume -M|--persistent-mount <VMFS UUID|label>` command to remount the vmfs volumes persistently with the same signature if the original 7-Mode LUNs are unmapped from the host or if the host is not online.

```
# esxcfg-volume -M vmfs-span-ds
```

1.6.3 Remediation of Raw Devices Mapped LUNs (RDM) attached to VMs


VMs with RDMs attached (pointing to a Data ONTAP LUN that was migrated) will fail to power on after transition. This is a result of the LUN serial number change that occurs when you transition from Data ONTAP operating in 7-Mode to clustered Data ONTAP.

You resolve this by removing and reattaching RDM disks to all VMs with RDMs based on the LUN serial number text/string mapping list (7-Mode LUN to clustered Data ONTAP LUN detail mapping) provided by 7MTT Tool.

1. In the Inventory Assessment Workbook, navigate to the "Host VM Disk Details" tab and identify the ESXi host with VM that has the "Type" column value of either PTRDM or NPTRDM. Note the VM name, and disk path details from "Disk" Column, and the naa id from the "Device Mapped" column.
2. Verify that the naa ID noted above is listed in the 7MTT mapping file that was generated post transition. Verify that the naa ID noted down has a corresponding new naa ID in the "LUN WWID" column of the mapping file. This will be the new clustered Data ONTAP LUN naa ID. Note the new clustered Data ONTAP LUN naa ID.
3. Use the clustered Data ONTAP LUN naa ID from the "LUN WWID" column that you captured from the mapping file, and the disk path details you captured from the Inventory Assessment Workbook, to reattach the clustered Data ONTAP LUN to the VM. For more information, see the section on reattaching RDMs.

1.6.3.1 Removing stale RDM Using vSphere Client:

1. Use the VM name and disk path of the RDM that you noted from the Inventory Assessment Workbook to detach it from the VM.
2. Open ESXi host or vCenter server managing the ESXi host using vSphere Client.
3. Right-click the VM and select *edit settings*. The VM Properties window is displayed.
4. Select the hard disk listed in the devices list using the disk path noted earlier from the Inventory Assessment Workbook.




5. Make note of the **Virtual Device Node** and **Compatibility Mode** from VM Properties window


For example:

Virtual Device Node: SCSI (0:2)

Compatibility Mode: Physical




6. Click Remove, and then select **Remove from virtual machine and delete files from disk** from the right pane.
7. Click OK. An error message similar to the following will appear. You can ignore this message and click Close.




This removes the stale RDM from the VM.

Reattaching RDM using vSphere Client:

1. Right-click the VM again and select *Edit Settings*. The VM Properties window opens.
2. Click *Add*. The Add Hardware window opens. Click **Hard Disk** and then click *Next* to select the disk.



3. Choose **Raw Device Mappings**. Click *Next* to select target LUN.



4. Select the LUN with the new clustered Data ONTAP NAA ID that you noted from the 7MTT mapping file, and then click *Next* to select **Datstore**.

Add Hardware

Select and Configure a Raw LUN
Which LUN would you like to use for this raw disk?

[Device Type](#)
[Select a Disk](#)
Select Target LUN
Select Datastore
Compatibility Mode
Advanced Options
Ready to Complete

Name, Identifier, Path ID, LUN or Capacity contain... Clear

Name	Identifier	Path ID	LUN
NETAPP...	naa.600a098054314c6c442b446f79716533	vmhba33:C0:T0:L...	36
NETAPP...	naa.600a098054314c6c442b446f79716534	vmhba33:C0:T0:L...	37
NETAPP...	naa.600a098054314c6c442b446f79716535	vmhba33:C0:T0:L...	38
NETAPP...	naa.600a098054314c6c442b446f79716536	vmhba33:C0:T0:L...	39
NETAPP...	naa.600a098054314c6c442b446f79716537	vmhba33:C0:T0:L...	40
NETAPP...	naa.600a098054314c6c4f5d447030754a78	vmhba33:C0:T0:L...	41
NETAPP...	naa.600a098054314c6c442b446f79716538	vmhba33:C0:T0:L...	42
NETAPP...	naa.600a098054314c6c4f5d447030754a79	vmhba33:C0:T0:L...	43
NETAPP...	naa.600a098054314c6c442b446f79716539	vmhba33:C0:T0:L...	44
NETAPP...	naa.600a098054314c6c442b446f7971652d	vmhba33:C0:T0:L...	45
NETAPP...	naa.600a098054314c6c442b446f79716541	vmhba33:C0:T0:L...	46
NETAPP...	naa.600a098054314c6c442b446f79716542	vmhba33:C0:T0:L...	47
NETAPP...	naa.600a098054314c6c442b446f79716543	vmhba33:C0:T0:L...	48
NETAPP...	naa.600a098054314c6c442b446f79716544	vmhba33:C0:T0:L...	49
NETAPP...	naa.600a098054314c6c442b446f79716545	vmhba33:C0:T0:L...	50
NETAPP...	naa.600a098054314c6c442b446f79716546	vmhba33:C0:T0:L...	51

Help ≤ Back Next ≥ Cancel

5. Choose **Select Datastore**. Select the datastore that matches the disk path that you noted in the 7MTT mapping file, and then click *Next*.

Add Hardware

Select a Datastore
Onto which datastore do you want to map this LUN?

[Device Type](#)
[Select a Disk](#)
[Select Target LUN](#)
Select Datastore
[Compatibility Mode](#)
[Advanced Options](#)
[Ready to Complete](#)


Select the datastore on which to store the LUN mapping. You will use the disk map on this datastore to access the virtual disk.

☒ Store with Virtual Machine
☐ Specify datastore


Datastore	# Hosts	Datastore Cluster
datastore1 (2)	1	N/A
vmfs3 (1)	2	N/A
vmfs5 (1)	2	N/A

[Help](#) [< Back](#) [Next >](#) [Cancel](#)

6. Choose the **Compatibility Mode** option, either Physical or Virtual. Compatibility Mode: Physical is selected in the following illustration. Click *Next*.



- Choose Advance Options, and then select **Virtual Device Node**, as indicated in the following illustration (SCSI (0:2)). Click Next.



- In the summary page, click *Finish* to submit your changes. Repeat the steps for all VMs with RDM attached.

Reattaching RDM using ESXi CLI/Console:

- Login to the ESXi console using SSH.
- Make a backup of the RDM disk file you noted from the “Disk” column of the Inventory Assessment Workbook. Use the mv command to back up the disk file and the associated device file.

For example:

If **/vmfs/volumes/53a3ac3d-df5aca03-3a94-001ec9d631cb/VM2-win-bus-A/VM2-win-bus-A.vmdk** was the RDM disk file you noted earlier:

```
# mv /vmfs/volumes/53a3ac3d-df5aca03-3a94-001ec9d631cb/VM2-win-bus-A/VM2-win-bus-A.vmdk /vmfs/volumes/53a3ac3d-df5aca03-3a94-001ec9d631cb/VM2-win-bus-A/VM2-win-bus-A.vmdk _bak
```

For Physical compatibility RDM (PTRDM), insert **-rdmp** into the file name:

```
# mv /vmfs/volumes/53a3ac3d-df5aca03-3a94-001ec9d631cb/VM2-win-bus-A/VM2-win-bus-A-rdmp.vmdk /vmfs/volumes/53a3ac3d-df5aca03-3a94-001ec9d631cb/VM2-win-bus-A/VM2-win-bus-A-rdmp.vmdk _bak
```

For Virtual compatibility RDM (NPTRDM), insert **-rdmp** into the file name:

```
# mv /vmfs/volumes/53a3ac3d-df5aca03-3a94-001ec9d631cb/VM2-win-bus-A/VM2-win-bus-A-rdm.vmdk /vmfs/volumes/53a3ac3d-df5aca03-3a94-001ec9d631cb/VM2-win-bus-A/VM2-win-bus-A-rdm.vmdk _bak
```

3. Re-create the RDM configuration and device files using the new clustered Data ONTAP LUN naa ID you noted from the 7MTT mapping file and the RDM disk file noted above from Inventory Assessment Workbook in the “Disk” column.

For Physical compatibility RDM (PTRDM), use the following command:

```
# vmkfstools -z /vmfs/devices/disks/<new clustered Data ONTAP naa id> <.vmdk configuration file>
```

For example:

```
# vmkfstools -z /vmfs/devices/disks/naa.600a098054314c6c442b446f79712313 /vmfs/volumes/53a3ac3d-df5aca03-3a94-001ec9d631cb/VM2-win-bus-A/VM2-win-bus-A.vmdk
```

For Virtual compatibility RDM (NPTRDM), use the following command:

```
# vmkfstools -r /vmfs/devices/disks/< new clustered Data ONTAP naa id > <.vmdk configuration file>
```

For example:

```
# vmkfstools -r /vmfs/devices/disks/naa.600a098054314c6c442b446f79712313 /vmfs/volumes/53a3ac3d-df5aca03-3a94-001ec9d631cb/VM2-win-bus-A/VM2-win-bus-A.vmdk
```

4. Execute below command to confirm that the configuration and pointer files are created:

```
# ls /vmfs/volumes/<datastore>/<VM directory>
#ls /vmfs/volumes/53a3ac3d-df5aca03-3a94-001ec9d631cb/VM2-win-bus-A
```

The new configuration and pointer files should be displayed under the mentioned VM directory path.

For example:

```
VM2-win-bus-A.vmdk
VM2-win-bus-A-rdm.vmdk
```

Repeat the steps for all VMs with RDM attached.

5. After remediation of all VMs with RDM in a given ESXi host, restart hostd and vpxa agent in the ESXi host by executing the following commands:

```
/etc/init.d/hostd restart
/etc/init.d/vpxa restart
```

1.6.4 Enabling Compare and Write (CAW) on ESX

By default, clustered Data ONTAP 8.3 and later supports CAW, also known as ATS in ESXi, which is supported only with newly created VMFS5 datastores. ESXi uses two settings to support this feature:

ATS global setting in ESXi:

ATS global setting is enabled by default in ESXi 5.x.

Complete the following steps to verify the global setting using vSphere ESXi 5.x CLI:

1. Check if VAAI features are enabled on the ESXi 5.x host from a console or SSH session using the following command:

```
# esxcli system settings advanced list -o /VMFS3/HardwareAcceleratedLocking
```

The output looks similar to this:

```
Path: /VMFS3/HardwareAcceleratedLocking
Type: integer
Int Value: 0
Default Int Value: 1
Min Value: 0
Max Value: 1
String Value:
Default String Value:
Valid Characters:
Description: Enable hardware accelerated VMFS locking (requires compliant hardware)
```

The Int Value of 0 indicates that VAAI ATS locking is disabled. Set the Int Value to 1 to enable ATS globally in ESXi host.

2. Enable ATS globally in ESXi by entering the following command:

```
# esxcli system settings advanced set -i 1 -o /VMFS3/HardwareAcceleratedLocking
```

For more information on enabling the HardwareAcceleratedLocking VAAI setting, see the VMware KB [1021976](#) - Frequently Asked Questions for vStorage APIs for Array Integration.

ATS per LUN/VMFS5 setting in ESXi:

The LUN level CAW/ATS setting is maintained in the VMFS5 header, which is set specifically for a given VMFS5 datastore at the time it is created depending on the LUN support for CAW from storage side. i.e while during the creation of VMFS5 datastore if the LUN and Storage backing it had CAW supported/enabled, then the VMFS header will be marked with ATS support. Similarly if LUN and Storage backing the VMFS5 datastore did not have CAW supported while creation of the datastore, then VMFS header will be marked with ATS not supported.

ATS per LUN/VMFS5 setting in ESXi must be enabled manually when migrating from Data ONTAP operating in 7-Mode without CAW support to Clustered Data ONTAP 8.3 and later with default CAW support. CAW is recommended by NetApp. See [KB 3013935](#).

Note: This task must be performed only during scheduled maintenance; for example, when no I/O or VMs are actively running on the VMFS datastore.

- If a datastore was migrated, ensure that the datastore is first remediated as described in the section above.
- In the Inventory Assessment Workbook that is generated by ICT, navigate

to the "SAN Host Filesystems" tab and verify CAW status for the datastore from the "ATS/CAW" column.

- If the "ATS/CAW" value for a datastore displays Disabled and if the "Filesystem" column is VMFS5.x then note the "Datastore Name", and complete the following steps to enable CAW on that datastore.

Enabling CAW on a datastore using ESXi CLI/Console

1. Login to ESXi console using SSH.
2. Using the above noted "Datastore Name" from Inventory Assessment Workbook, execute the following command to list device and partition details:

```
~ # vmkfstools -Ph -vl <datastore path>
```

For example:

```
~ # vmkfstools -Ph -vl /vmfs/volumes/datastorename
```

```
VMFS-5.60 file system spanning 1 partitions.
File system label (if any): datastorename
Mode: public
Capacity 9.8 GB, 8.2 GB available, file block size 1 MB, max file size 64
TB
Volume Creation Time: Mon Dec 9 10:29:18 2013
Files (max/free): 27408/27394
Ptr Blocks (max/free): 64512/64495
Sub Blocks (max/free): 3968/3964
Secondary Ptr Blocks (max/free): 256/256
File Blocks (overcommit/used/overcommit %): 0/1593/0
Ptr Blocks (overcommit/used/overcommit %): 0/17/0
Sub Blocks (overcommit/used/overcommit %): 0/4/0
Volume Metadata size: 590675968
UUID: 52a59b7e-52d2fb6c-11d6-001ec9d631cb
Partitions spanned (on "lvm"):
  naa.600a098044314c6c442b446d51376749:1
  naa.600a098054314c6c445d446f79716431:1
  naa.600a098054314c6c445d446f79716433:1
Is Native Snapshot Capable: YES
```

3. Note the first device name along with partition number.
For instance **naa.600a098044314c6c442b446d51376749:1** from the previous example.
4. Use this device ID along with the partition number to enable CAW on the datastore using the following command:

```
~# vmkfstools --configATSONly 1 /vmfs/devices/disks/device-ID:Partition
```

For example:

```
~# vmkfstools --configATSONly 1 /vmfs/devices/disks/naa.600a098044314c6c442b446d51376749:1
```

Refer to [this VMWare document](#) for more information.

5. To verify that the VMFS volume has been configured with ATS only, enter the following command:

```
# vmkfstools -Ph -vl /vmfs/volumes/VMFS-volume-name
```

```
VMFS-5.54 file system spanning 1 partitions.
```

```
File system label (if any): ats-test-1
```

```
Mode: public ATS-only
```

Guest OS post remediation procedures:

Guest operating systems on the VMs might require additional post remediation steps after transition of LUN's to clustered Data ONTAP.

LINUX

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

WINDOWS

1. Using Disk Manager, bring online the disks.
2. If Windows Failover Cluster is configured on the VM, bring online the disks from Failover Cluster Manager.

1.7 Configure recommended settings for ESXi host

After post remediation of the ESXi host, you must ensure that the recommended ESXi host settings for clustered Data ONTAP are applied for the ESXi hosts. ESXi host settings can be configured on the host by using **Virtual Storage Console (VSC)**. VSC is the standard NetApp plugin that enables vSphere vCenter to configure ESXi host settings for Data ONTAP. ESXi hosts and VMs deployed on the source 7-Mode system should be configured using VSC. Alternatively, you can configure VMs manually by using the information in the following KB articles:

- [Guest OS tunings](#)
- [Task Set Full \(QFull\) Tunables for LUNs in vSphere 5.1](#)
- [Storage Array Type Plugin option for a NetApp array on VMware vSphere](#)
- [HardwareAcceleratedLocking setting required for VMware deployment](#)

7MTT Linux Host Remediation Steps

RHEL 5 series

Data LUN Scenarios

1.8 DMMP device without file system

This scenario describes a LUN configured with device mapper multipath (DMMP) and without a file system. A scenario using “alias” name configured on the DMMP devices is also included.

1.8.1 Pretransition

1. To identify the LUN name to be transitioned, see the “LUN Name” column in the Inventory Assessment Workbook’s “LUN” tab.
2. Get the SCSI device names for the transition LUN from “SCSI Device ID” column of the “SAN Host LUNs” tab.

Alternatively, you can locate the SCSI device name for a LUN on host by using the “`sanlun lun show`” command.

For example, when the transition LUN is “dmmp_raw”, the respective SCSI devices are “/dev/sdo and /dev/sdk”.

When the transition LUN is “dmmp_raw_alias”, the respective SCSI devices are “/dev/sdp and /dev/ sdl”

```
[root@IBMX3550M3-229-169 ~]# sanlun lun show
```

controller(7mode)/ vserver(Cmode)	lun-pathname	device filename	host adapter	protocol	lun size	mode
fas3270-229-147	/vol/gm_dmmp_raw_vol/dmmp_raw	/dev/sdo	host12	iSCSI	1g	7
fas3270-229-147	/vol/gm_dmmp_raw_alias_vol/dmmp_raw_alias	/dev/sdp	host12	iSCSI	1g	7
fas3270-229-147	/vol/gm_dmmp_raw_vol/dmmp_raw	/dev/sdk	host11	iSCSI	1g	7
fas3270-229-147	/vol/gm_dmmp_raw_alias_vol/dmmp_raw_alias	/dev/sdl	host11	iSCSI	1g	7

3. From the Inventory Assessment Workbook’s “SAN Host LUNs” tab, identify the DMMP device name for the transition LUN using the column titled “OS Device ID”. Alternatively, you can identify the DMMP device name for the LUN using the “`multipath -ll`” command.

The DMMP device name can be a device handle ID (WWID) or alias. For example, you might have two DMMP devices: one device has an alias configured and the other device uses device handle ID, or ‘naa’ ID.

The DMMP device name for SCSI IDs “/dev/sdo and /dev/sdk” is
“360a980003753456258244538554b4b53”.

The DMMP alias name for SCSI IDs “/dev/sdp and /dev/ sdl” is “dmmp_raw_lun”.

```
[root@IBMX3550M3-229-169 ~]# multipath -ll
dmmp_raw_lun (360a980003753456258244538554b4b55) dm-8 NETAPP,LUN
[size=1.0G][features=3 queue_if_no_path pg_init_retries 50][hwhandler=0][rw]
\_ round-robin 0 [prio=2][active]
\_ 11:0:0:2 sdl 8:176 [active][ready]
\_ 12:0:0:2 sdp 8:240 [active][ready]

360a980003753456258244538554b4b53 dm-7 NETAPP,LUN
[size=1.0G][features=3 queue_if_no_path pg_init_retries 50][hwhandler=0][rw]
\_ round-robin 0 [prio=2][active]
\_ 11:0:0:1 sdk 8:160 [active][ready]
\_ 12:0:0:1 sdo 8:224 [active][ready]
```

4. Go to “SAN Host File system” tab to ensure that the identified LUN for transition does not have a file system. If the DMMP device entry is not present then a file system is not configured on that device. Alternatively, you can use the “dumpe2fs” command to ensure that the identified LUN for transition does not have a file system configured.

Example:

This example illustrates the error message displayed when the “dumpe2fs” command is executed against the DMMP device name and alias name where the file system is not created. You can determine from the output that the DMMP device does not have a file system, and is, therefore, a RAW device.

```
[root@IBMX3550M3-229-169 ~]# dumpe2fs /dev/mapper/360a980003753456258244538554b4b53
dumpe2fs 1.39 (29-May-2006)
dumpe2fs: Bad magic number in super-block while trying to open /dev/mapper/360a980003753456258244538554b4b53
Couldn't find valid filesystem superblock.

[root@IBMX3550M3-229-169 ~]# dumpe2fs /dev/mapper/dmmp_raw_lun
dumpe2fs 1.39 (29-May-2006)
dumpe2fs: Bad magic number in super-block while trying to open /dev/mapper/dmmp_raw_lun
Couldn't find valid filesystem superblock.
```

1.8.2 Before entering cutover phase

For FC configurations, ensure that fabric connectivity and zoning to clustered Data ONTAP controllers exists. When you use iSCSI configurations, ensure that iSCSI sessions to clustered Data ONTAP controllers are discovered and logged in.

1. Stop I/O to the mount points.
2. Ensure the applications that are accessing the LUNs are gracefully shut down.
3. Flush the 7-Mode LUN's DMMP device or aliases using the “multipath -f <device_name>” command. See the pretransition steps for locating the DMMP device or aliases for the respective LUNs. You can also refer to the Inventory Assessment Workbook’s “SAN Host LUNs” tab and the “OS Device ID” column to get DMMP device names.

Example:

```
[root@IBMX3550M3-229-169 ~]# multipath -f 360a980003753456258244538554b4b53

[root@IBMX3550M3-229-169 ~]# multipath -f dmmp_raw_lun
```

1.8.3 Posttransition

Posttransition steps begin on the host after the LUNs are transitioned to clustered Data ONTAP. After the transition, the 7-Mode volumes are in the “offline” state, and the 7-Mode LUNs are not accessible to the host.

After LUN transition, the DMMP device handle ID (WWID) changes. You must either manually replace the 7-Mode LUN’s WWID with the clustered Data ONTAP LUN’s WWID across all necessary files and applications, or you must create a multipath alias with the 7-Mode LUN’s WWID. Both methods are described in this section.

NOTE: After the transition completes, the 7MTT Tool provides a LUN mapping table between 7-Mode LUNs and their corresponding clustered Data ONTAP LUNs. You can use this table to identify the clustered Data ONTAP LUN device handle IDs and LUN path names.

1. Remove the SCSI devices that were created for 7-Mode LUNs. See the pretransition steps to get the SCSI device IDs. You use the “rescan-scsi-bus.sh -r” command to remove all SCSI devices in the host using a single command. You might also decide to delete each SCSI device individually using the “echo 1 > /sys/block/<scsi_id>/device/delete” command. The latter command must be executed on all 7-Mode LUN SCSI devices. Refer to the Inventory Assessment Workbook’s “SAN Host LUNs” tab and the “SCSI Device ID” column to identify the LUN’s SCSI device IDs.

Example:

Delete SCSI devices “/dev/sdo and /dev/sdk” for DMMP device
360a980003753456258244538554b4b53

Delete SCSI devices “/dev/sdp and /dev/ sdl” for DMMP alias dmmp_raw_lun.

```
[root@IBMX3550M3-229-169 /]# echo 1 > /sys/block/sdk/device/delete
[root@IBMX3550M3-229-169 /]# echo 1 > /sys/block/sdo/device/delete

[root@IBMX3550M3-229-169 /]# echo 1 > /sys/block/sdl/device/delete
[root@IBMX3550M3-229-169 /]# echo 1 > /sys/block/sdp/device/delete
```

2. Execute “rescan-scsi-bus.sh” command to rescan new clustered Data ONTAP LUNs.
3. Use the “sanlun lun show” command to obtain the new SCSI device name for the clustered Data ONTAP LUN.

Example:

The clustered Data ONTAP LUN’s “dmmp_raw” SCSI devices display as “/dev/sdac, /dev/sdr, /dev/sdz and /dev/sds”.

```
[root@IBMX3550M3-229-169 ~]# sanlun lun show
```

vserver(Cmode)	lun-pathname	filename	adapter	protocol	size	mode
vs_7MTT	/vol/gm_cdot_dmmp_raw_vol/dmmp_raw	/dev/sdac	host7	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_dmmp_raw_vol/dmmp_raw	/dev/sdr	host5	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_dmmp_raw_vol/dmmp_raw	/dev/sdz	host8	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_dmmp_raw_vol/dmmp_raw	/dev/sds	host6	iSCSI	1g	C

The clustered Data ONTAP LUN’s “dmmp_raw_alias” SCSI devices display as “/dev/sdab, /dev/sdaa, /dev/sdy and /dev/ sdz”.

```
[root@IBMX3550M3-229-169 ~]# sanlun lun show
```

vserver(Cmode)	lun-pathname	filename	adapter	protocol	size	mode
vs_7MTT	/vol/gm_cdot_dmmp_raw_alias_vol/dmmp_raw_alias	/dev/sdab	host8	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_dmmp_raw_alias_vol/dmmp_raw_alias	/dev/sdaa	host7	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_dmmp_raw_alias_vol/dmmp_raw_alias	/dev/sdy	host6	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_dmmp_raw_alias_vol/dmmp_raw_alias	/dev/sdx	host5	iSCSI	1g	C

The device handle ID (WWID) for the clustered Data ONTAP LUNs can be obtained using the command “/sbin/scsi_id -g -u -s /block/<SCSI dev>”.

Example:

The WWID for LUN “dmmp_raw” is “3600a098051764b2d4f3f453135452d32”

The WWID for LUN “dmmp_raw_alias” is “3600a098051764b2d4f3f453135452d31”

```
[root@IBMX3550M3-229-169 ~]# scsi_id -g -u -s /block/sdac
3600a098051764b2d4f3f453135452d32
```

```
[root@IBMX3550M3-229-169 ~]# scsi_id -g -u -s /block/sdab
3600a098051764b2d4f3f453135452d31
```

1.8.3.1 Method 1: Find and Replace

1. In a RAW LUN scenario, you must replace the 7-Mode device handle ID with the clustered Data ONTAP LUN’s device handle ID across all the applicable configuration files, such as /etc/fstab, and across applications, for example, Oracle DB.
2. In scenarios where DMMP raw device is configured with an DMMP alias, you must update the “/etc/multipath.conf” file to replace the 7-Mode LUN device handle ID with the clustered Data ONTAP LUN device handle ID. To get the new clustered Data ONTAP LUN device handle ID, refer to the “LUN mapping table” that is generated by the 7MTT tool.

Example:

In this example, the 7-Mode device handle ID “360a980003753456258244538554b4b55” (LUN name - dmmp_raw_alias) has a multipath alias “dmmp_raw_lun”. Replace the device handle ID with the clustered Data ONTAP LUN ID “3600a098051764b2d4f3f453135452d31”.

Example output before the update:

```
[root@IBMX3550M3-229-169 ~]# cat /etc/multipath.conf
defaults {
    user_friendly_names      no
    queue_without_daemon     no
    flush_on_last_del        yes
    max_fds                  max
    pg_prio_calc              avg
}
blacklist {
    wwid                     "3600605b0041c58d016fe219e0fd0ee92"
    devnode                  "^cciss.*"
}
multipaths {
    multipath {
        wwid                 360a980003753456258244538554b4b55
        alias                 dmmp_raw_lun
    }
}
```

Example output after the update:

```
[root@IBMX3550M3-229-169 ~]# cat /etc/multipath.conf
defaults {
    user_friendly_names          no
    queue_without_daemon         no
    flush_on_last_del            yes
    max_fds                      max
    pg_prio_calc                 avg
}
blacklist {
    wwid                        "3600605b0041c58d016fe219e0fd0ee92"
    devnode                     "^cciss.*"
}
multipaths {
    multipath {
        wwid                    3600a098051764b2d4f3f453135452d31
        alias                    dmmp_raw_lun
    }
}
```

3. Execute the “multipath” command to update the DMMP alias settings.
4. Verify that the DMMP alias name correctly references the clustered Data ONTAP LUN ID.

Example:

In the sample output, the DMMP alias “dmmp_raw_lun” points to ID “3600a098051764b2d4f3f453135452d31”

```
[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][hw_handler=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]
```

1.8.3.2 Method 2: Create new multipath alias for DMMP raw LUN

1. If you do not want to manually replace device handle IDs, you can instead create a new multipath alias configuration, where a new alias name with the 7-Mode LUN's device handle ID is created to reference the clustered Data ONTAP LUN's device handle ID. A new DMMP alias is created by editing the /etc/multipath.conf file and adding new a “multipath” subsection after the “multipaths” section.

Example:

This example displays a portion of the /etc/multipath.conf file with a new alias entry after the “multipaths” section. The new alias name is “360a980003753456258244538554b4b53” designating the 7-Mode LUN's device handle ID. This alias name references the new clustered Data ONTAP LUN's device handle ID: “3600a098051764b2d4f3f453135452d32”.

```
multipaths {
    multipath {
        wwid                    3600a098051764b2d4f3f453135452d32
        alias                    360a980003753456258244538554b4b53
    }
}
```

2. Verify that the DMMP alias is correctly referencing the clustered Data ONTAP LUN ID.

Example:

The DMMP alias “360a980003753456258244538554b4b53” points to ID “3600a098051764b2d4f3f453135452d32”.

```
[root@IBMX3550M3-229-169 ~]# multipath -ll 360a980003753456258244538554b4b53
360a980003753456258244538554b4b53 (3600a098051764b2d4f3f453135452d32) dm-7 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]
\_ 5:0:0:5 sdr 65:16 [active][ready]
\_ 8:0:0:5 sdz 65:144 [active][ready]
\_ round-robin 0 [prio=10][enabled]
\_ 7:0:0:5 sdac 65:192 [active][ready]
\_ 6:0:0:5 sds 65:32 [active][ready]
```

File system on DMMP device

1.9 Mount point using DMMP device name

1.9.1 Pretransition

1. To identify the LUN name that is being transitioned, refer to the Inventory Assessment Workbook’s “LUN Name” column in the “LUN” tab.
2. Locate the SCSI device names for the transition LUN from the “SCSI Device ID” column in the “SAN Host LUNs” tab. Alternatively, you can locate the SCSI device name for a given LUN on a host by running the “sanlun lun show” command.

Example:

The transition LUN is “dmmp_fs” and the respective SCSI devices are “/dev/sdg and /dev/sdc”.

```
[root@IBMX3550M3-229-169 ~]# sanlun lun show
```

vserver (Cmode)	lun-pathname	filename	adapter	protocol	size	mode
fas3270-229-146	/vol/gm_dmmp_fs_vol/dmmp_fs	/dev/sdg	host10	iSCSI	1g	7
fas3270-229-146	/vol/gm_dmmp_fs_vol/dmmp_fs	/dev/sdc	host9	iSCSI	1g	7

3. You can identify the DMMP device name for the transition LUN in the “OS Device ID” column in the Inventory Assessment Workbook’s “SAN Host LUNs” tab. Alternatively, you can identify the DMMP device name for the LUN using the “multipath -ll” command.

Example:

The DMMP device name for SCSI IDs “/dev/sdc and /dev/sdg” is “360a9800037534562572b453855496b41”.

```
[root@IBMX3550M3-229-169 ~]# multipath -ll /dev/sdc
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]
\_ 10:0:0:1 sdg 8:96 [active][ready]
```

4. You can identify the file system configured on the DMMP device in the “Filesystems” column of the

Inventory Assessment Workbook's "SAN Host Filesystems" tab. Alternatively, you can identify the file system configured on DMMP device using the "blkid" command. Check the TYPE value against the DMMP device.

Example:

The file system created on DMMP device "360a9800037534562572b453855496b41" is "ext3", as shown in the sample output:

```
[root@IBMX3550M3-229-169 ~]# blkid | grep -i 360a9800037534562572b453855496b41
/dev/mapper/360a9800037534562572b453855496b41: UUID="a073547e-00b6-4bf9-8e08-5eef08499a9c" SEC_TYPE="ext2" TYPE="ext3"
```

5. You can identify the file system UUID number for the LUN in the "UUID" column in the Inventory Assessment Workbook's "Host Filesystems" tab. Alternatively, you can use the "dumpe2fs <device_path_name> | grep UUID" command to identify the file system UUID number for the LUN.

Example:

The UUID for DMMP device "360a9800037534562572b453855496b41" is "a073547e-00b6-4bf9-8e08-5eef08499a9c", as shown in the sample output:

```
[root@IBMX3550M3-229-169 ~]# dumpe2fs /dev/mapper/360a9800037534562572b453855496b41 | grep UUID
dumpe2fs 1.39 (29-May-2006)
Filesystem UUID:          a073547e-00b6-4bf9-8e08-5eef08499a9c
```

6. You can identify the directory on which the DMMP device is mounted by looking in the Inventory Assessment Workbook's "Mount" column in the "SAN Hosts Filesystems" tab. Alternatively, you can enter either the "df -h" or the "mount" command to find where the DMMP device is mounted.

Example:

Using the "df -h" command, the sample output indicates that DMMP device "/dev/mapper/360a9800037534562572b453855496b41" is mounted on directory "/mnt/dmmp_ext3":

```
[root@IBMX3550M3-229-169 ~]# df -h
Filesystem                                Size  Used Avail Use% Mounted on
/dev/mapper/360a9800037534562572b453855496b41
1008M   34M  924M   4% /mnt/dmmp_ext3
```

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

Example:

The sample output indicates that DMMP device

"/dev/mapper/360a9800037534562572b453855496b41" is present in "/etc/fstab" file:

```
[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/LogVol01 swap swap defaults 0 0
/dev/mapper/test_vg-test_lv /mnt/lvm_ext3 ext3 defaults,_netdev 0 0
/dev/mapper/360a9800037534562572b453855496b41 /mnt/dmmp_ext3 ext3 defaults,_netdev 0 0
```

8. Create a backup of the “/etc/fstab” file before editing the file

Example:

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

9. Edit the “/etc/fstab” file and replace the DMMP device name with its respective file system UUID number.

Example:

In this 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
```

1.9.2 Before entering cutover phase

1. For FC configurations, ensure that fabric connectivity and zoning to clustered Data ONTAP controllers is established. When you use iSCSI configurations, ensure that iSCSI sessions to clustered Data ONTAP controllers are discovered and logged in.
2. Ensure that the applications that access the LUNs are shut down gracefully.
3. Stop I/O to the mount points.
4. Unmount DMMP devices using the “umount <dir_name>” command.

Example:

```
# umount /mnt/dmmp_ext3
```

5. Flush the 7-Mode LUN's DMMP device ID using the “multipath -f <device_name>” command. See the pretransition steps to identify the DMMP device or alias for respective 7-Mode LUNs. In the Inventory Assessment Workbook, you can also refer to the “OS Device ID” column in the “SAN Host LUNs” tab to identify the DMMP device names.

Example:

```
[root@IBMX3550M3-229-169 ~]# multipath -f 360a9800037534562572b453855496b41
```

1.9.3 Posttransition

Posttransition steps begin on the host after the LUNs are transitioned from 7-Mode to clustered Data ONTAP. After the transition, the 7-Mode volumes are in the “offline” state, and the 7-Mode LUNs are not accessible to

the host.

Also after LUN transition, the DMMP device handle ID (WWID) changes.

NOTE: After the transition is complete, the 7MTT Tool provides a LUN mapping table between 7-Mode LUNs and the corresponding clustered Data ONTAP LUNs. You can use this table to identify the clustered Data ONTAP LUN device handle IDs and LUN path names.

1. Remove the SCSI devices that were created for 7-Mode LUNs. See the pretransition steps to get the SCSI device IDs. To remove all SCSI devices in the host with a single command, use the “`rescan-scsi-bus.sh -r`” command. You can also delete each SCSI device individually by using the “`echo 1 > /sys/block/<scsi_id>/device/delete`” command on all 7-Mode LUN SCSI devices. To locate the new clustered Data ONTAP LUN device handle ID, refer to the “LUN mapping table” generated by the 7MTT tool.

Example:

Delete SCSI devices “/dev/sdj and /dev/sdl” for DMMP device
360a9800037534562572b453855496b41

```
[root@IBMx3550M3-229-169 ~]# echo 1 > /sys/block/sdj/device/delete
[root@IBMx3550M3-229-169 ~]# echo 1 > /sys/block/sdl/device/delete
```

2. To discover new clustered Data ONTAP LUNs, use the “`rescan-scsi-bus.sh`” command.
3. Verify clustered Data ONTAP LUNs using “`sanlun lun show`”

Example:

In the sample output, the clustered Data ONTAP LUN’s “dmmp_fs” SCSI devices are “/dev/sdq, /dev/sdo, /dev/sdm and /dev/sdk”:

```
[root@IBMx3550M3-229-169 ~]# sanlun lun show
```

vserver(Cmode)	lun-pathname	filename	adapter	protocol	size	mode
vs_7MTT	/vol/gm_cdot_dmmp_fs_vol/dmmp_fs	/dev/sdq	host8	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_dmmp_fs_vol/dmmp_fs	/dev/sdo	host7	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_dmmp_fs_vol/dmmp_fs	/dev/sdm	host6	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_dmmp_fs_vol/dmmp_fs	/dev/sdk	host5	iSCSI	1g	C

4. Use the “`multipath`” command to configure DMMP devices for clustered Data ONTAP LUNs.
5. Use the “`multipath -ll`” command to check the device handle ID for clustered Data ONTAP LUNs.

Example:

The clustered Data ONTAP LUN’s “dmmp_fs” DMMP device handle ID is
“3600a098051764937303f4479515a7451”

```
[root@IBMX3550M3-229-169 ~]# multipath -ll /dev/sdq
3600a098051764937303f4479515a7451 dm-6 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][active]
\_ 6:0:0:4 sdm 8:192 [active][ready]
\_ 7:0:0:4 sdo 8:224 [active][ready]
\_ round-robin 0 [prio=10][enabled]
\_ 5:0:0:4 sdk 8:160 [active][ready]
\_ 8:0:0:4 sdq 65:0 [active][ready]
```

6. If the mount points are defined in the “/etc/fstab” file, use the “mount -a” command to mount the LUN. Since the device handle ID was replaced by its corresponding file system UUID in the pretransition steps, there is no need to update the “/etc/fstab” file with clustered Data ONTAP LUN’s device handle ID.
7. Using the “mount” command, verify that the DMMP device is mounted.

Example:

In the sample output, the clustered Data ONTAP LUN

“/dev/mapper/3600a098051764937303f4479515a7451” is not mounted on directory “/mnt/dmmp_ext3”.

```
[root@IBMX3550M3-229-169 ~]# mount
/dev/mapper/VolGroup00-LogVol00 on / type ext3 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
/dev/sda1 on /boot type ext3 (rw)
tmpfs on /dev/shm type tmpfs (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
/dev/mapper/test_vg-test_lv on /mnt/lvm_ext3 type ext3 (rw)
/dev/mapper/3600a098051764937303f4479515a7451 on /mnt/dmmp_ext3 type ext3 (rw,_netdev)
```

1.10 Mount point using alias name

1.10.1 Pretransition

1. To identify the LUN being transitioned, see the Inventory Assessment Workbook’s “LUN Name” column in the “LUN” tab.
2. To identify the SCSI device names for the transition LUN, see the “SCSI Device ID” column in the Inventory Assessment Workbook’s “SAN Host LUNs” tab. Alternatively, you can identify the SCSI device name for a LUN using the “sanlun lun show” command.

Example:

In the sample output, the transition LUN is “dmmp_fs_alias” and the respective SCSI devices are “/dev/sdd and /dev/sdh”:

```
[root@IBMX3550M3-229-169 ~]# sanlun lun show
```

vserver(Cmode)	lun-pathname	filename	adapter	protocol	size	mode
fas3270-229-146	/vol/gm_dmmp_fs_alias_vol/dmmp_fs_alias	/dev/sdh	host10	iSCSI	1g	7
fas3270-229-146	/vol/gm_dmmp_fs_alias_vol/dmmp_fs_alias	/dev/sdd	host9	iSCSI	1g	7

3. To identify the DMMP device name for the transition LUN, see the “OS Device ID” column in the

“SAN Host LUNs” tab. Alternatively, you can identify the DMMP device name for the LUN by using the “multipath -ll” command.

Example:

In the sample output, the alias for SCSI IDs “/dev/sdd and /dev/ sdd” is “dmmp_fs_lun”.

```
[root@IBMX3550M3-229-169 ~]# multipath -ll /dev/sdd
dmmp_fs_lun (360a9800037534562572b453855496b43) dm-4 NETAPP,LUN
[size=1.0G][features=3 queue_if_no_path pg_init_retries 50][hwhandler=0][rw]
\_ round-robin 0 [prio=2][active]
\_ 9:0:0:2 sdd 8:48 [active][ready]
\_ 10:0:0:2 sdh 8:112 [active][ready]
```

4. From the Inventory Assessment Workbook’s “SAN Host Filesystems” tab, identify the file system configured on the DMMP device. Alternatively, you can identify the file system using the “blkid” command, and then checking the TYPE value against that of the DMMP device.

Example:

In the sample output, the file system created on the DMMP alias “dmmp_fs_lun” is “ext3”:

```
[root@IBMX3550M3-229-169 ~]# blkid | grep -i dmmp_fs_lun
/dev/mapper/dmmp_fs_lun: UUID="690321df-06be-4750-9262-9e533fe8bbfd" SEC_TYPE="ext2" TYPE="ext3"
```

5. From the “UUID” column of the Inventory Assessment Workbook’s “SAN Hosts Filesystems” tab, identify the file system “UUID” number for the LUN. Alternatively, you can use the “dumpe2fs <device_path_name> | grep UUID” command to identify the file system.

Example:

In the sample output, the UUID for the DMMP alias “dmmp_fs_lun” is “690321df-06be-4750-9262-9e533fe8bbfd”:

```
[root@IBMX3550M3-229-169 ~]# dumpe2fs /dev/mapper/dmmp_fs_lun | grep UUID
dumpe2fs 1.39 (29-May-2006)
Filesystem UUID: 690321df-06be-4750-9262-9e533fe8bbfd
```

6. From the “Mount” column of the Inventory Assessment Workbook’s “SAN Hosts Filesystems” tab, identify the directory on which the DMMP device is mounted. Alternatively, you can use either the “df -h” command or the “mount” command to identify the directory.

Example:

The DMMP alias device “/dev/mapper/dmmp_fs_lun” is mounted on directory “/mnt/dmmp_alias_ext3”

```
[root@IBMX3550M3-229-169 ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/mapper/dmmp_fs_lun
                1008M   34M   924M   4% /mnt/dmmp_alias_ext3
```

1.10.2 Before entering cutover phase

1. For FC configurations, ensure that fabric connectivity and zoning to clustered Data ONTAP controllers is established. When you use iSCSI configurations, you must ensure that iSCSI sessions to clustered Data ONTAP controllers are discovered and logged in.
2. Stop I/O to the mount points.
3. Ensure that applications that are accessing the LUNs are shut down gracefully.
4. Unmount the DMMP device using the “`umount <dir_name>`” command.

Example:

```
# umount /mnt/dmmp_alias_ext3/
```

5. Flush the 7-Mode LUN's DMMP aliases using the “`multipath -f <device_name>`” command. See the pretransition steps for identifying the DMMP device or aliases for the respective 7-Mode LUNs. You can also go to the Inventory Assessment Workbook's “SAN Host LUNs” tab to get the DMMP device names.

Example:

```
[root@IBMX3550M3-229-169 ~]# multipath -f dmmp_fs_lun
```

1.10.3 Posttransition

Posttransition steps begin on the host after the LUNs are transitioned to clustered Data ONTAP. After the transition, the 7-Mode volumes are in the “offline” state, and the 7-Mode LUNs are not accessible to the host.

After LUN transition the DMMP alias name does not change; however, the DMMP device handle ID (WWID) does change. You must modify the “`/etc/multipath.conf`” file to reference the alias to the clustered Data ONTAP LUN's device handle ID.

NOTE: After the transition completes, the 7MTT Tool provides a LUN mapping table between 7-Mode LUNs and the corresponding clustered Data ONTAP LUNs. You can use this table to identify clustered Data ONTAP LUN's device handle IDs and LUN path names.

1. Remove the SCSI devices that were created for 7-Mode LUNs. See the pretransition steps for identifying the SCSI device IDs. You can use either the “`rescan-scsi-bus.sh -r`” command to remove all SCSI devices in the host at once, or you can delete each SCSI device individually using the “`echo 1 > /sys/block/<scsi_id>/device/delete`” command on all 7-Mode LUN SCSI devices. To identify a LUN's SCSI devices, see the “SCSI Device ID” column in the Inventory Assessment Workbook's “SAN Host LUNs” tab.

Example:

The following example deletes SCSI devices “`/dev/sdk` and “`/dev/sdm`” for the DMMP alias `dmmp_fs_lun`:

```
[root@IBMX3550M3-229-169 ~]# echo 1 > /sys/block/sdk/device/delete
[root@IBMX3550M3-229-169 ~]# echo 1 > /sys/block/sdm/device/delete
```

2. Use the “`rescan-scsi-bus.sh`” command to discover new clustered Data ONTAP LUNs.
3. Verify the clustered Data ONTAP LUNs using “`sanlun lun show`”.

Example:

In the sample output, the clustered Data ONTAP LUN's "dmmp_fs_alias" SCSI devices are "/dev/sdp, /dev/sdn, /dev/sdl & /dev/sdj":

```
[root@IBMx3550M3-229-169 ~]# sanlun lun show
```

vserver(Cmode)	lun-pathname	filename	adapter	protocol	size	mode
vs_7MTT	/vol/gm_cdot_dmmp_fs_alias_vol/dmmp_fs_alias	/dev/sdp	host8	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_dmmp_fs_alias_vol/dmmp_fs_alias	/dev/sdn	host7	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_dmmp_fs_alias_vol/dmmp_fs_alias	/dev/sdl	host6	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_dmmp_fs_alias_vol/dmmp_fs_alias	/dev/sdj	host5	iSCSI	1g	C

4. In the "/etc/multipath.conf" file, replace the 7-Mode device handle ID with the clustered Data ONTAP LUN's device handle ID below the "multipaths" section of the file so that the "alias name" points to the clustered Data ONTAP LUN ID. To locate the clustered Data ONTAP device handle ID, see the "LUN mapping table" that is generated by the 7MTT tool.

Example:

The sample output before the update displays the 7-Mode LUN ID

"360a9800037534562572b453855496b43" pointing to the alias "dmmp_fs_lun":

```
multipaths {
    multipath {
        wwid      360a980003753456258244538554b4b55
        alias     dmmp_raw_lun
    }
    multipath {
        wwid      360a9800037534562572b453855496b43
        alias     dmmp_fs_lun
    }
}
```

After the update, the corresponding clustered Data ONTAP LUN ID,

"3600a098051764937303f4479515a7452", points to the alias "dmmp_fs_lun":

```
multipaths {
    multipath {
        wwid      360a980003753456258244538554b4b55
        alias     dmmp_raw_lun
    }
    multipath {
        wwid      3600a098051764937303f4479515a7452
        alias     dmmp_fs_lun
    }
}
```

5. Use the "multipath" command to configure the DMMP devices for clustered Data ONTAP LUNs.
6. Verify that the DMMP alias now points to the clustered Data ONTAP LUN device handle ID.

Example:

```
[root@IBMX3550M3-229-169 ~]# multipath -ll /dev/sdn
dmmp_fs_lun (3600a098051764937303f4479515a7452) dm-5 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]
\_ 6:0:0:2 sdl 8:176 [active][ready]
\_ 7:0:0:2 sdn 8:208 [active][ready]
\_ round-robin 0 [prio=10][enabled]
\_ 5:0:0:2 sdj 8:144 [active][ready]
\_ 8:0:0:2 sdp 8:240 [active][ready]
```

7. Mount the clustered Data ONTAP LUN to its respective mount point directory using the alias in the command, enter the “mount /dev/mapper/<alias_name> <mount_dir_name>” command.

Example:

```
# mount /dev/mapper/dmmp_fs_lun /mnt/dmmp_alias_ext3/
```

8. If the mount points are defined in the “/etc/fstab” file, use “mount -a” command to mount the LUN.
9. Use the “mount” command to verify that the DMMP device is mounted.

Example:

In the sample output, the clustered Data ONTAP LUN “/dev/mapper/dmmp_fs_lun” is not mounted on the “/mnt/dmmp_alias_ext3” directory:

```
[root@IBMX3550M3-229-169 ~]# mount
/dev/mapper/VolGroup00-LogVol00 on / type ext3 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
/dev/sda1 on /boot type ext3 (rw)
tmpfs on /dev/shm type tmpfs (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
/dev/mapper/test_vg-test_lv on /mnt/lvm_ext3 type ext3 (rw)
/dev/mapper/3600a098051764937303f4479515a7451 on /mnt/dmmp_ext3 type ext3 (rw,_netdev)
/dev/mapper/dmmp_fs_lun on /mnt/dmmp_alias_ext3 type ext3 (rw,_netdev)
```

1.11 File system on LVM device

1.11.1 Pretransition

1. From the “LUN Name” column of the Inventory Assessment Workbook’s “LUN” tab, identify the name of the LUN being transitioned.
2. From the “SCSI DeviceID” column of the “SAN Host LUNs” tab, locate the SCSI device names for the transition LUN. Alternatively, you can identify the SCSI device name for a LUN by using the “sanlun lun show” command.

Example:

In the sample output, the transition LUNs are “lvm1, lvm2, lvm3” and their respective SCSI devices are “/dev/sdn, /dev/sdj, /dev/sdf, /dev/sdb, /dev/sdi, /dev/sde”:

```
[root@IBMX3550M3-229-169 ~]# sanlun lun show
```

controller(7mode)/ vserver(Cmode)	lun-pathname	device filename	host adapter	protocol	lun size	mode
fas3270-229-147	/vol/gm_lvm1_vol/lvm1	/dev/sdn	host12	iSCSI	1g	7
fas3270-229-147	/vol/gm_lvm1_vol/lvm1	/dev/sdj	host11	iSCSI	1g	7
fas3270-229-146	/vol/gm_lvm3_vol/lvm3	/dev/sdi	host10	iSCSI	1g	7
fas3270-229-146	/vol/gm_lvm2_vol/lvm2	/dev/sdf	host10	iSCSI	1g	7
fas3270-229-146	/vol/gm_lvm3_vol/lvm3	/dev/sde	host9	iSCSI	1g	7
fas3270-229-146	/vol/gm_lvm2_vol/lvm2	/dev/sdb	host9	iSCSI	1g	7

- From the “OS Device ID” column of the Inventory Assessment Workbook’s “SAN Host LUNs” tab, identify the DMMP device name for the transition LUN. Alternatively, you can identify the DMMP device handle ID for the LUN using “multipath -ll” command. The DMMP device name can be a device handle ID (WWID) or alias name

Example:

This following example depicts three DMMP devices that are used for LVM configuration.

The DMMP device name for SCSI IDs “/dev/sdj and /dev/sdn” is

“360a980003753456258244538554b4b51”.

The DMMP device name for SCSI IDs “/dev/sdb and /dev/sdf” is

“360a9800037534562572b453855496b37”.

The DMMP device name for SCSI IDs “/dev/sde and /dev/sdi” is

“360a9800037534562572b453855496b39”.

```
[root@IBMX3550M3-229-169 ~]# multipath -ll /dev/sdn
360a980003753456258244538554b4b51 dm-6 NETAPP,LUN
[size=1.0G][features=3 queue_if_no_path pg_init_retries 50][hwhandler=0][rw]
\_ round-robin 0 [prio=2][active]
\_ 11:0:0:0 sdj 8:144 [active][ready]
\_ 12:0:0:0 sdn 8:208 [active][ready]

[root@IBMX3550M3-229-169 ~]# multipath -ll /dev/sdb
360a9800037534562572b453855496b37 dm-2 NETAPP,LUN
[size=1.0G][features=3 queue_if_no_path pg_init_retries 50][hwhandler=0][rw]
\_ round-robin 0 [prio=2][active]
\_ 9:0:0:0 sdb 8:16 [active][ready]
\_ 10:0:0:0 sdf 8:80 [active][ready]

[root@IBMX3550M3-229-169 ~]# multipath -ll /dev/sde
360a9800037534562572b453855496b39 dm-5 NETAPP,LUN
[size=1.0G][features=3 queue_if_no_path pg_init_retries 50][hwhandler=0][rw]
\_ round-robin 0 [prio=2][active]
\_ 9:0:0:3 sde 8:64 [active][ready]
\_ 10:0:0:3 sdi 8:128 [active][ready]
```

- From the “Physical Volume Name” column of the Inventory Assessment Workbook’s SAN Host LVMs” tab, identify the DMMP devices that are used by LVM. Alternatively, you can use the “pvscan” command to identify the DMMP devices that are used by LVM.

Example:

The following example shows three DMMP devices being used in an LVM configuration:

```
[root@IBMX3550M3-229-169 ~]# pvscan
PV /dev/mpath/360a9800037534562572b453855496b37 VG test_vg lvm2 [1020.00 MB / 0 free]
PV /dev/mpath/360a980003753456258244538554b4b51 VG test_vg lvm2 [1020.00 MB / 0 free]
PV /dev/mpath/360a9800037534562572b453855496b39 VG test_vg lvm2 [1020.00 MB / 1012.00 MB free]
```

- From the “Volume Group Name” and “Logical Volume Path” columns in the Inventory Assessment Workbook’s “SAN Host LVMs” tab, identify the Volume Group (VG) and Logical Volume (LV) names respectively. Alternatively, you can use the “vgscan” and “lvscan” commands respectively to determine the VG and LV names.

Example:

In this example, the VG name is “test_vg” and the LV name is “test_lv”:

```
[root@IBMX3550M3-229-169 ~]# vgscan
Reading all physical volumes. This may take a while...
Found volume group "test_vg" using metadata type lvm2

[root@IBMX3550M3-229-169 ~]# lvscan
ACTIVE                '/dev/test_vg/test_lv' [2.00 GB] inherit
```

- You can display verbose information of VG and LV by using the “vgdisplay” and “lvdisplay” commands respectively.
- From the “Filesystem” column of the Inventory Assessment Workbook’s “SAN Host Filesystems” tab, identify the file system configured on the LV device. Alternatively, you can identify the file system configured on the LV device by using the “blkid” command and then checking the TYPE value against that of the LV device.

Example:

In the sample output, the file system created on Logical Volume “test_lv” is “ext3”:

```
[root@IBMX3550M3-229-169 ~]# blkid | grep -i test_lv
/dev/test_vg/test_lv: UUID="86da0704-d328-45b4-95dd-d485a0aa5c43" SEC_TYPE="ext2" TYPE="ext3"
/dev/mapper/test_vg-test_lv: UUID="86da0704-d328-45b4-95dd-d485a0aa5c43" SEC_TYPE="ext2" TYPE="ext3"
```

- From the “Mount” column of the Inventory Assessment Workbook’s “SAN Host Filesystems” tab, identify the directory on which the LVs are mounted. Alternatively, you can use either of the “df -h” command or the “mount” command.

Example:

In the sample output, the LV “/dev/mapper/test_vg-test_lv” is mounted on the directory “/mnt/lvm_ext3”:

```
[root@IBMX3550M3-229-169 ~]# df -h
Filesystem              Size  Used Avail Use% Mounted on
/dev/mapper/test_vg-test_lv
                        2.0G   68M   1.9G   4% /mnt/lvm_ext3
```

1.11.2 Before entering cutover phase

1. For FC configurations, ensure that fabric connectivity and zoning to clustered Data ONTAP controllers is established. If you are using iSCSI configurations, you must ensure that iSCSI sessions to clustered Data ONTAP controllers are discovered and logged in.
2. Stop I/O to LV mount points.
3. Ensure that applications that are accessing the LUNs are shut down gracefully.
4. Unmount the LV mount point using the “`umount <dir_name>`” command.

Example:

```
# umount /mnt/lvm_ext3
```

5. Disable Logical Volume using the “`vgchange -an <vg_name>`” command.

Example:

```
[root@IBMX3550M3-229-169 ~]# vgchange -an test_vg
0 logical volume(s) in volume group "test_vg" now active
```

6. Verify Logical Volume (LV) status using the “`lvdisplay`” command. The LV status displays as “NOT available”.

Example:

```
[root@IBMX3550M3-229-169 ~]# lvdisplay test_vg
--- Logical volume ---
LV Name                /dev/test_vg/test_lv
VG Name                test_vg
LV UUID                qUydTw-y8xX-hp3H-hHDg-FGgB-a3WT-E6sG61
LV Write Access        read/write
LV Status              NOT available
LV Size                2.00 GB
Current LE             512
Segments              3
Allocation             inherit
Read ahead sectors    auto
```

7. Export Volume Group using the “`vgexport <vg_name>`” command.

Example:

```
[root@IBMX3550M3-229-169 ~]# vgexport test_vg
Volume group "test_vg" successfully exported
```

8. Verify Volume Group (VG) status using the “`vgdisplay`” command. The VG status display as “exported”.

Example:

```
[root@IBMx3550M3-229-169 ~]# lvdisplay test_vg
Volume group test_vg is exported

[root@IBMx3550M3-229-169 ~]# vgdisplay test_vg
Volume group test_vg is exported
--- Volume group ---
VG Name                test_vg
System ID
Format                 lvm2
Metadata Areas         3
Metadata Sequence No   3
VG Access              read/write
VG Status              exported/resizable
MAX LV                 0
Cur LV                1
Open LV                0
Max PV                 0
Cur PV                3
Act PV                 3
VG Size                2.99 GB
PE Size                4.00 MB
Total PE               765
Alloc PE / Size        512 / 2.00 GB
Free PE / Size         253 / 1012.00 MB
VG UUID                UL36f0-iOxk-31qh-0lHa-0Lso-Ljkm-5gsgBB
```

9. Flush the 7-Mode LUN's DMMP device IDs using the "multipath -f <device_name>" command. See the pretransition steps to get the DMMP device or alias names for the respective 7-Mode LUNs.

Example:

```
[root@IBMx3550M3-229-169 ~]# multipath -f 360a9800037534562572b453855496b37

[root@IBMx3550M3-229-169 ~]# multipath -f 360a9800037534562572b453855496b39

[root@IBMx3550M3-229-169 ~]# multipath -f 360a980003753456258244538554b4b51
```

1.11.3 Posttransition

Posttransition steps begin on the host after the LUNs are transitioned to clustered Data ONTAP. After the transition, the 7-Mode volumes are in the "offline" state, and the 7-Mode LUNs are not accessible to the host.

After LUN transition, the LVM attributes, for example, the LV/VG names, do not change. Therefore, the same LV and VG names are used during posttransition configuration.

NOTE: After the transition completes, 7MTT Tool provides a LUN mapping table between 7-Mode LUNs and their corresponding clustered Data ONTAP LUNs. You can use this table to identify the clustered Data ONTAP LUN device handle IDs and LUN path names.

1. Remove the SCSI devices that were created for 7-Mode LUNs. See the pretransition steps for

verifying the SCSI device IDs. To remove all SCSI devices in the host, use the “rescan-scsi-bus.sh -r” command. To individually delete each SCSI device, use the “echo 1 > /sys/block/<scsi_id>/device/delete” command on all 7-Mode LUN SCSI devices. To identify a LUN’s SCSI device ID, see the “SCSI Device ID” column in the Inventory Assessment Workbook’s “SAN Host LUNs” tab.

Example:

Delete SCSI devices “/dev/sdj & /dev/sdn” for DMMP device
360a980003753456258244538554b4b51

Delete SCSI devices “/dev/sdb & /dev/sdf” for DMMP device
360a9800037534562572b453855496b37

Delete SCSI devices “/dev/sde & /dev/sdi” for DMMP device
360a9800037534562572b453855496b39

```
[root@IBMX3550M3-229-169 ~]# echo 1 > /sys/block/sdj/device/delete
[root@IBMX3550M3-229-169 ~]# echo 1 > /sys/block/sdn/device/delete

[root@IBMX3550M3-229-169 ~]# echo 1 > /sys/block/sdb/device/delete
[root@IBMX3550M3-229-169 ~]# echo 1 > /sys/block/sdf/device/delete

[root@IBMX3550M3-229-169 ~]# echo 1 > /sys/block/sde/device/delete
[root@IBMX3550M3-229-169 ~]# echo 1 > /sys/block/sdi/device/delete
```

2. Execute the “rescan-scsi-bus.sh” command to discover new clustered Data ONTAP LUNs.
3. Execute the “multipath” command to configure DMMP devices for clustered Data ONTAP LUNs.

Example:

```
[root@IBMX3550M3-229-169 ~]# multipath
reload: 3600a098051764937303f4479515a7450 NETAPP,LUN C-Mode
[size=1.0G][features=3 queue_if_no_path pg_init_retries 50][hwhandler=0][n/a]
\_ round-robin 0 [prio=4][undef]
\_ 14:0:0:0 sdj 8:144 [active][ready]
\_ 15:0:0:0 sdr 65:16 [active][ready]
\_ round-robin 0 [prio=1][undef]
\_ 13:0:0:0 sdf 8:80 [active][ready]
\_ 16:0:0:0 sdt 65:48 [active][ready]
reload: 3600a098051764937303f4479515a744f NETAPP,LUN C-Mode
[size=1.0G][features=3 queue_if_no_path pg_init_retries 50][hwhandler=0][n/a]
\_ round-robin 0 [prio=4][undef]
\_ 14:0:0:3 sdn 8:208 [active][ready]
\_ 15:0:0:3 sds 65:32 [active][ready]
\_ round-robin 0 [prio=1][undef]
\_ 13:0:0:3 sdi 8:128 [active][ready]
\_ 16:0:0:3 sdu 65:64 [active][ready]
```

4. Verify the clustered Data ONTAP LUNs using the “sanlun lun show” command.

Example:

In the sample output, the clustered Data ONTAP LUNs “lvm1” SCSI devices are “/dev/sdw, /dev/sdv, /dev/sdu and /dev/sdt”, the clustered Data ONTAP LUNs “lvm2” SCSI devices are “/dev/sde, /dev/sdd, /dev/sdc and /dev/sdb”, and the clustered Data ONTAP LUNs “lvm3” SCSI devices are “/dev/sdi, /dev/sdf, /dev/sdh

and /dev/sdg”:

```
[root@IBMX3550M3-229-169 ~]# sanlun lun show
```

vserver (Cmode)	lun-pathname	filename	adapter	protocol	size	mode
vs_7MTT	/vol/gm_cdot_lvm1_vol/lvm1	/dev/sdw	host8	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_lvm1_vol/lvm1	/dev/sdv	host7	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_lvm1_vol/lvm1	/dev/sdu	host6	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_lvm1_vol/lvm1	/dev/sdt	host5	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_lvm3_vol/lvm3	/dev/sdi	host8	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_lvm3_vol/lvm3	/dev/sdf	host5	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_lvm3_vol/lvm3	/dev/sdh	host7	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_lvm3_vol/lvm3	/dev/sdg	host6	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_lvm2_vol/lvm2	/dev/sde	host8	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_lvm2_vol/lvm2	/dev/sdd	host7	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_lvm2_vol/lvm2	/dev/sdc	host6	iSCSI	1g	C
vs_7MTT	/vol/gm_cdot_lvm2_vol/lvm2	/dev/sdb	host5	iSCSI	1g	C

5. To determine the new clustered Data ONTAP LUN device handle ID, see the “LUN mapping table” that is generated by the 7MTT tool. Alternatively, you can use the the “multipath -ll” command to identify the clustered Data ONTAP LUN’s DMMP device handle ID.

Example:

In the sample output, the clustered Data ONTAP LUN “lvm1” device ID is 3600a098051764b2d4f3f453135452d30, the clustered Data ONTAP LUN “lvm2” device ID is 3600a098051764937303f4479515a7450, and the clustered Data ONTAP LUN “lvm3” device ID is 3600a098051764937303f4479515a744f:

```
[root@IBMX3550M3-229-169 ~]# multipath -ll /dev/sdw
3600a098051764b2d4f3f453135452d30 dm-9 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][active]
\_ 5:0:0:1 sdt 65:48 [active][ready]
\_ 8:0:0:1 sdw 65:96 [active][ready]
\_ round-robin 0 [prio=10][enabled]
\_ 6:0:0:1 sdu 65:64 [active][ready]
\_ 7:0:0:1 sdv 65:80 [active][ready]

[root@IBMX3550M3-229-169 ~]# multipath -ll /dev/sdb
3600a098051764937303f4479515a7450 dm-2 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][active]
\_ 6:0:0:0 sdc 8:32 [active][ready]
\_ 7:0:0:0 sdd 8:48 [active][ready]
\_ round-robin 0 [prio=10][enabled]
\_ 5:0:0:0 sdb 8:16 [active][ready]
\_ 8:0:0:0 sde 8:64 [active][ready]

[root@IBMX3550M3-229-169 ~]# multipath -ll /dev/sdi
3600a098051764937303f4479515a744f dm-3 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][active]
\_ 6:0:0:3 sdg 8:96 [active][ready]
\_ 7:0:0:3 sdh 8:112 [active][ready]
\_ round-robin 0 [prio=10][enabled]
\_ 5:0:0:3 sdf 8:80 [active][ready]
\_ 8:0:0:3 sdi 8:128 [active][ready]
```

6. Import Volume Group using the “vgimport <vg_name>” command.

Example:

```
[root@IBMX3550M3-229-169 ~]# vgimport test_vg
Volume group "test_vg" successfully imported
```

7. Verify Volume Group (VG) status using the “vgdisplay” command.

Example:

The output of “vgdisplay” indicates that the VG is no longer exported:

```
[root@IBMX3550M3-229-169 ~]# vgdisplay test_vg
--- Volume group ---
VG Name                test_vg
System ID
Format                 lvm2
Metadata Areas         3
Metadata Sequence No   6
VG Access               read/write
VG Status               resizable
MAX LV                 0
Cur LV                 1
Open LV                 0
Max PV                 0
Cur PV                 3
Act PV                 3
VG Size                 2.99 GB
PE Size                 4.00 MB
Total PE                765
Alloc PE / Size         512 / 2.00 GB
Free PE / Size           253 / 1012.00 MB
VG UUID                UL36f0-iOxk-31qh-0lHa-0Lso-LjkM-5gsgBB
```

8. Enable Logical Volumes (LV) using the “vgchange -an <vg_name>” command.

Example:

```
[root@IBMX3550M3-229-169 ~]# vgchange -ay test_vg
1 logical volume(s) in volume group "test_vg" now active
```

9. Verify Logical Volume (LV) status using the “lvdisplay” command. The LV status displays as “available”.

Example:

```
[root@IBMX3550M3-229-169 ~]# lvdisplay test_vg
--- Logical volume ---
LV Name                /dev/test_vg/test_lv
VG Name                test_vg
LV UUID                qUydTw-ySxX-hp3H-hHDg-FGgB-a3WT-E6sG61
LV Write Access        read/write
LV Status              available
# open                 0
LV Size                2.00 GB
Current LE             512
Segments               3
Allocation             inherit
Read ahead sectors     auto
- currently set to    256
Block device           253:4
```

10. Physical Volume (PV) status can be identified using the “pvdisplay” command or the “pvscan” command.

Example:

The “pvscan” command output indicates that the VG “test_vg” has 3 DMMP devices:

```
[root@IBMX3550M3-229-169 ~]# pvscan
PV /dev/mpath/3600a098051764937303f4479515a7450 VG test_vg lvm2 [1020.00 MB / 0 free]
PV /dev/mpath/3600a098051764b2d4f3f453135452d30 VG test_vg lvm2 [1020.00 MB / 0 free]
PV /dev/mpath/3600a098051764937303f4479515a744f VG test_vg lvm2 [1020.00 MB / 1012.00 MB free]
```

11. Mount the LVs from the clustered Data ONTAP LUN to its respective mount point directory using the “mount <lv_name> <mount_dir_name>” command.

Example:

```
# mount /dev/test_vg/test_lv /mnt/lvm_ext3
```

12. If the mount points are defined in the “/etc/fstab” file, use the “mount -a” command to mount the LVs.
13. Verify the mount points using the “mount” command.

Example:

```
[root@IBMX3550M3-229-169 ~]# mount
/dev/mapper/VolGroup00-LogVol100 on / type ext3 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
/dev/sda1 on /boot type ext3 (rw)
tmpfs on /dev/shm type tmpfs (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
/dev/mapper/test_vg-test_lv on /mnt/lvm_ext3 type ext3 (rw)
```

SAN Boot LUN Scenarios

1.12 Protocol: FC/FCoE

1.12.1 Pretransition

1. From the 7-Mode controller's console, identify the 7-Mode LUN name on which RHEL 5 OS is installed. See the "LUN Name" column of the Inventory Assessment Workbook's "LUN" tab. Alternatively, on the host, use the 7-Mode command "lun show" to list the LUNs on controller.
2. From the "SCSI Device ID" column of the Inventory Assessment Workbook's "SAN Host LUNs" tab, identify the SCSI device names for the transition LUN from the. Alternatively, you can get the SCSI device name on a host for a given LUN by using the "sanlun lun show" command.

Example:

The transition LUN is "san_lun_5" and the respective SCSI devices are "/dev/sda, /dev/sdb & /dev/sdc"

```
[root@IBMX3550-229-5 ~]# sanlun lun show
controller(7mode)/
vserver(Cmode)      lun-pathname      device
                    filename      host
                    adapter  protocol  lun
                    size    mode
-----
fas3270-229-146     /vol/rhel5_fc_sanboot/san_lun_5 /dev/sda      host0      FCP      55.0g  7
fas3270-229-146     /vol/rhel5_fc_sanboot/san_lun_5 /dev/sdb      host0      FCP      55.0g  7
fas3270-229-146     /vol/rhel5_fc_sanboot/san_lun_5 /dev/sdc      host1      FCP      55.0g  7
```

3. From the OS Device ID column of the "SAN Host LUNs" tab, identify the DMMP device name for the transition LUN. Alternatively, you can use the "multipath -ll" command to identify the DMMP device name for the LUN on the host.

Example:

DMMP device name for SCSI IDs "/dev/sda, /dev/sdb & /dev/sdc" is "360a9800037534562572b453855496b4f"

```
[root@IBMX3550-229-5 ~]# multipath -ll /dev/sda
360a9800037534562572b453855496b4f dm-0 NETAPP,LUN
[size=55G][features=3 queue_if_no_path pg_init_retries 50][hwhandler=0][rw]
\_ round-robin 0 [prio=4][active]
\_ 1:0:4:0 sdc 8:32 [active][ready]
\_ round-robin 0 [prio=1][enabled]
\_ 0:0:4:0 sda 8:0 [active][ready]
\_ 0:0:5:0 sdb 8:16 [active][ready]
```

4. Check that the DMMP device exists in the "/dev/mapper" directory by running the "ls /dev/mapper/ <DMMP device name>" command. If you cannot locate the device, then the DMMP device might be using an alias or user-friendly name.

Example:

In this example, "ls /dev/mapper/360a9800037534562572b453855496b4f", and DMMP devices is not using either "user_friendly_name" or "alias".

```
[root@IBMX3550-229-5 ~]# ls /dev/mapper/360a9800037534562572b453855496b4f
/dev/mapper/360a9800037534562572b453855496b4f
```

5. Identify the DMMP devices and LVM names on which RHEL 5 OS's "/boot" and "/" (root) partitions are installed. See the "SCSI Device ID / Device Name" and "Mount" columns of the Inventory Assessment Workbook's "SAN Host Filesystems" tab to identify the mount directory.

Alternatively, from the host, you can use either the "df -h" or "mount" command.

Example:

The "/boot" partition is created on the DMMP device "/dev/mapper/ and then 360a9800037534562572b453855496b4fp1" is mounted on the directory "/boot". The "p1" designation represents the partition on the DMMP device.
The "/" (root) partition is installed on Logical Volume "/dev/mapper/VolGroup00-LogVol00".

```
[root@IBMx3550-229-5 ~]# df -h
Filesystem                Size      Used Avail Use% Mounted on
/dev/mapper/VolGroup00-LogVol00
                          27G       3.1G   23G   13% /
/dev/mapper/360a9800037534562572b453855496b4fp1
                          99M        27M   68M   29% /boot
tmpfs                     20G         0    20G    0% /dev/shm
```

```
[root@IBMx3550-229-5 ~]# mount
/dev/mapper/VolGroup00-LogVol00 on / type ext3 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
/dev/mapper/360a9800037534562572b453855496b4fp1 on /boot type ext3 (rw)
tmpfs on /dev/shm type tmpfs (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
```

6. From the "Filesystem" column of the Inventory Assessment Workbook's "SAN Host Filesystems" tab, identify the file system configured on the DMMP device.

Alternatively, you can use the "blkid" command to identify the file system, and check the TYPE value against that of the DMMP and LV devices.

Example:

The file system created on DMMP device (/boot) "360a9800037534562572b453855496b4f" and LV (/) "/dev/mapper/VolGroup00-LogVol00" is "ext3":

```
[root@IBMx3550-229-5 ~]# blkid
/dev/mapper/VolGroup00-LogVol01: TYPE="swap"
/dev/VolGroup00/LogVol01: TYPE="swap"
/dev/mapper/VolGroup00-LogVol00: UUID="d91ef637-43bf-42a7-a7c8-9c1e5329d17a" TYPE="ext3"
/dev/sda1: LABEL="/boot" UUID="5e842e74-caa2-451f-bce4-d2d2e302b1e0" TYPE="ext3"
/dev/sdb1: LABEL="/boot" UUID="5e842e74-caa2-451f-bce4-d2d2e302b1e0" TYPE="ext3"
/dev/sdc1: LABEL="/boot" UUID="5e842e74-caa2-451f-bce4-d2d2e302b1e0" TYPE="ext3"
/dev/mapper/360a9800037534562572b453855496b4fp1: LABEL="/boot" UUID="5e842e74-caa2-451f-bce4-d2d2e302b1e0" TYPE="ext3"
```

7. If the "/" (root) partition is installed on Logical Volume (LV), then no pretransition changes to the configuration are required.

Note: By default, RHEL 5 installs the "/" partition on LV.

8. If the "/boot" partition is installed on DMMP device, confirm how it is referenced to mount in "/etc/fstab" at boot time.

Example:

The “/boot” partition is reference by its DMMP device name

“/dev/mapper/360a9800037534562572b453855496b4fp1”, and this device name is used at boot time to mount the “/boot” partition.

```
[root@IBMX3550-229-5 ~]# cat /etc/fstab
/dev/VolGroup00/LogVol00 / ext3 defaults 1 1
#/dev/mapper/mpath0p1 /boot ext3 defaults 1 2
/dev/mapper/360a9800037534562572b453855496b4fp1 /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/LogVol01 swap swap defaults 0 0
```

9. If the “/boot” partition is referenced by its DMMP device name then the DMMP device name should be replaced with file system UUID name in “/etc/fstab”. This step is required since the DMMP device name will change after the 7-Mode LUN is transition to clustered Data ONTAP.
10. From the “UUID” column of the Inventory Assessment Workbook’s “SAN Host Filesystem” tab, identify the file system “UUID” number of “/boot” partition.

Alternatively, you may identify the file system “UUID” by using “`dumpe2fs <device_path_name> | grep UUID`” command.

Example:

The UUID for “/boot” partition, for instance DMMP device

“360a9800037534562572b453855496b4fp1”, is “5e842e74-caa2-451f-bce4-d2d2e302b1e0”.

```
[root@IBMX3550-229-5 ~]# dumpe2fs /dev/mapper/360a9800037534562572b453855496b4fp1 | grep UUID
dumpe2fs 1.39 (29-May-2006)
Filesystem UUID: 5e842e74-caa2-451f-bce4-d2d2e302b1e0
```

11. Make a backup of “/etc/fstab” file before editing it.

Example:

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

12. Edit the “/etc/fstab” file and replace the DMMP device name with its respective file system UUID number.

Example:

In this example the “/boot” partition “/dev/mapper/360a9800037534562572b453855496b4fp1” is replaced by UUID “5e842e74-caa2-451f-bce4-d2d2e302b1e0”:

```
[root@IBMX3550-229-5 ~]# cat /etc/fstab
/dev/VolGroup00/LogVol00 / ext3 defaults 1 1
UUID=5e842e74-caa2-451f-bce4-d2d2e302b1e0 /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/LogVol01 swap swap defaults 0 0
```

13. Because the “/etc/fstab” entry for “/boot” partition has changed, the kernel initial RAM disk (initramfs) must be recreated. From the “Kernel” and “initrd” columns of the Inventory Assessment Workbook’s

“GRUB Configuration” tab, identify the initrd image to be modified.

Make a backup of the image file before recreating an initrd image.

Example:

The backup of initrd image is saved with the name “initrd-2.6.18-371.el5.img.bak”:

```
[root@IBMX3550-229-5 ~]# cp /boot/initrd-2.6.18-371.el5.img /boot/initrd-2.6.18-371.el5.img.bak
```

14. Obtain the information for the Grub “device.map” file from the “GRUB Device Map” tab in the Inventory Assessment Workbook.
15. Recreate the initrd-image using the command “mkinitrd -f /boot/ initrd-`uname -r`.img `uname -r` --with multipath”.

Example:

```
[root@IBMX3550-229-5 ~]# mkinitrd -f /boot/initrd-`uname -r`.img `uname -r`
```

16. Restart the host to boot from the new initrd image.

1.12.2 Before entering cutover phase

1. For FC configurations, ensure that fabric connectivity and zoning to clustered Data ONTAP controllers is established. For iSCSI configurations, ensure that iSCSI sessions to clustered Data ONTAP controllers are discovered and logged-in.
2. Shut down the host.

1.12.3 Posttransition

Post- transition steps on the host begin after the LUN transition to clustered Data ONTAP completes. After transition completes, the 7-Mode volumes will be in the “offline” state, so the 7-Mode LUNs are not accessible to the host.

After LUN transition the DMMP device handle ID (WWID) will change.

NOTE: After the transition completes, the 7MTT Tool provides a LUN mapping table between 7-Mode LUNs and their corresponding clustered Data ONTAP LUNs. You can use this table to identify the clustered Data ONTAP LUN device handle IDs and LUN path names.

1. Power-on the host.
2. Enter HBA BIOS settings mode.
3. Choose “Rescan” to discover clustered Data ONTAP SANBoot lun.
4. Clear the 7-Mode Boot LUN ID.
5. Add clustered Data ONTAP Boot LUN ID in HBA BIOS.
6. Exit HBA BIOS and reboot the host.
7. Wait for OS boot to complete.
8. Verify clustered Data ONTAP LUNs using “sanlun lun show”.
9. Execute “multipath -ll” to verify DMMP device.

Protocol: iSCSI

RHEL 5.x iSCSI SAN boot scenario is not supported. Therefore, the OS must be newly deployed on clustered Data ONTAP LUNs provisioned over iSCSI transport.

RHEL 6 series

Data LUN Scenarios

1.13 DMMP device without file system

1.13.1 Pretransition

1. From the “LUN Name” column of the Inventory Assessment Workbook’s “LUN” tab, identify the name of the LUN being transitioned.
2. From the “SCSI Device ID” column of the SAN Host LUNs” tab in the Inventory Assessment Workbook, identify the SCSI device names for the transition LUN. Alternatively, you can use the “sanlun lun show” command to identify the SCSI device name.

Example:

In the sample output, the transition LUN “lun_dmmp_raw” is associated with SCSI device “/dev/sdg” and the LUN “lun_dmmp_raw_with_alias” is associated with SCSI device “/dev/sdh”:

```
[root@IBMx3550M3-229-167 mapper]# sanlun lun show
```

controller (7mode)/ vservers (Cmode)	lun-pathname	device filename	host adapter	protocol	lun size	mode
fas3270-229-147	/vol/lun_dmmp_raw_with_alias_vol/lun_dmmp_raw_with_alias	/dev/sdh	host6	iSCSI	2g	7
fas3270-229-147	/vol/lun_dmmp_raw_vol/lun_dmmp_raw	/dev/sdg	host6	iSCSI	2g	7
fas3270-229-147	/vol/lun_lvm3_vol/lun_lvm3	/dev/sdf	host6	iSCSI	2g	7
fas3270-229-147	/vol/lun_lvm2_vol/lun_lvm2	/dev/sde	host6	iSCSI	2g	7
fas3270-229-146	/vol/lun_lvm1_vol/lun_lvm1	/dev/sdd	host5	iSCSI	2g	7
fas3270-229-146	/vol/lun_dmmp_fs_with_alias_vol/lun_dmmp_fs_with_alias	/dev/sdc	host5	iSCSI	2g	7
fas3270-229-146	/vol/lun_dmmp_fs_without_alias_vol/lun_dmmp_fs_without_alias	/dev/sdb	host5	iSCSI	2g	7

3. From the “OS Device ID” column of the Inventory Assessment Workbook’s “SAN Host LUNs” tab, identify the DMMP device name for the transition LUN.

Alternatively, you can use the “multipath -ll” command to identify the DMMP device name for the LUN. The DMMP device name is either a unique device handle ID (WWID) or an alias.

Example:

In the sample output, the DMMP device name for SCSI device “/dev/sdg” is “360a980003753456258244538554b4b4d”, and for “/dev/sdh”, the DMMP device is “raw_device”:

```
[root@IBMx3550M3-229-167 ~]# multipath -ll
360a9800037534562572b453855496b35 dm-5 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
+- policy='round-robin 0' prio=2 status=active
  - 5:0:0:2 sdd 8:48 active ready running
360a980003753456258244538554b4b4d dm-8 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
+- policy='round-robin 0' prio=2 status=active
  - 6:0:0:2 sdg 8:96 active ready running
360a9800037534562572b453855496b33 dm-4 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
+- policy='round-robin 0' prio=2 status=active
  - 5:0:0:1 sdc 8:32 active ready running
raw_device (360a980003753456258244538554b4b4f) dm-9 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
+- policy='round-robin 0' prio=2 status=active
  - 6:0:0:3 sdh 8:112 active ready running
```

4. Verify that the DMMP device exists in the “/dev/mapper” directory by running the “ls /dev/mapper/ <DMMP device name>” command. If the DMMP device is not displayed, the device might be using an alias or a user-friendly name.

Example:

For example, entering “ls /dev/mapper/raw_device” and “ls /dev/mapper/360a980003753456258244538554b4b4d” returns the DMMP device names, which indicates that the devices exist:

```
[root@IBMX3550M3-229-167 ~]# ls /dev/mapper/raw_device
/dev/mapper/raw_device
[root@IBMX3550M3-229-167 ~]# ls /dev/mapper/360a980003753456258244538554b4b4d
/dev/mapper/360a980003753456258244538554b4b4d
```

- From the Inventory Assessment Workbook’s “SAN Host LVMs” tab, ensure that the LUN identified for transition does not have LVM configured. If the LV device entry is not present then LVM is not configured on that device. Alternatively, you can determine whether the DMMP device is part of an LVM by running “blkid” command. Look at the TYPE value of the DMMP device. If the TYPE is “LVM2_member”, the DMMP is part of an LVM.

Example:

In the following example, the DMMP device “360a980003753456258244538554b4b4d” and “raw_device” are not part of an LVM:

```
[root@IBMX3550M3-229-167 mapper]# blkid
/dev/mapper/vg_ibmx3550m3229167-lv_root: UUID="4119733f-b52e-4bef-bbcb-2d2a623b6fdc" TYPE="ext4"
/dev/sda1: UUID="c875432c-1f86-4ac0-9cb7-5d832b23c377" TYPE="ext4"
/dev/sda2: UUID="hSq1V6-JrvF-Ppfi-kmYM-Haf0-RAx6-Q2I8Ho" TYPE="LVM2_member"
/dev/mapper/vg_ibmx3550m3229167-lv_swap: UUID="77014c28-2e68-4881-8b11-f24d4de58440" TYPE="swap"
/dev/mapper/vg_ibmx3550m3229167-lv_home: UUID="119e8728-d544-4572-a3f0-5c0c0ecf3e7f" TYPE="ext4"
/dev/mapper/360a9800037534562572b453855496b31: UUID="a21b6900-c8f8-4d5c-b97c-9528c999b3e7" TYPE="ext4"
/dev/mapper/360a9800037534562572b453855496b35: UUID="dLp0zP-PYCo-LiYl-W346-BZi8-zL5K-gQzKHd" TYPE="LVM2_member"
/dev/mapper/360a980003753456258244538554b4b4b: UUID="J2hkjU-0X5t-Ntgk-SkHb-3BW2-CqD9-0oRLUb" TYPE="LVM2_member"
/dev/mapper/360a980003753456258244538554b4b49: UUID="wKjYtR-mqwF-iTWg-sn9o-1v3Y-9Sy3-3lGutA" TYPE="LVM2_member"
/dev/mapper/n_vg-nvol: UUID="3c740d52-b8af-4ae9-b476-9754fac6748f" TYPE="ext4"
```

- From the Inventory Assessment Workbook’s “SAN Host File System” tab, ensure the LUN identified for transition does not have a file system configured. If the DMMP device entry is not present then the file system is not configured on that device.

Alternatively, you can use running the “blkid” command to determine whether the DMMP device has a file system. Check whether the TYPE value of the DMMP device indicates a file system exists.

Example:

In the following example, the DMMP devices “360a980003753456258244538554b4b4d” and “raw_device” do not have a file system:

```
[root@IBMX3550M3-229-167 mapper]# blkid
/dev/mapper/vg_ibmx3550m3229167-lv_root: UUID="4119733f-b52e-4bef-bbcb-2d2a623b6fdc" TYPE="ext4"
/dev/sda1: UUID="c875432c-1f86-4ac0-9cb7-5d832b23c377" TYPE="ext4"
/dev/sda2: UUID="hSq1V6-JrvF-Ppfi-kmYM-Haf0-RAx6-Q2I8Ho" TYPE="LVM2_member"
/dev/mapper/vg_ibmx3550m3229167-lv_swap: UUID="77014c28-2e68-4881-8b11-f24d4de58440" TYPE="swap"
/dev/mapper/vg_ibmx3550m3229167-lv_home: UUID="119e8728-d544-4572-a3f0-5c0c0ecf3e7f" TYPE="ext4"
/dev/mapper/360a9800037534562572b453855496b31: UUID="a21b6900-c8f8-4d5c-b97c-9528c999b3e7" TYPE="ext4"
/dev/mapper/360a9800037534562572b453855496b35: UUID="dLp0zP-PYCo-LiYl-W346-BZi8-zL5K-gQzKHd" TYPE="LVM2_member"
/dev/mapper/360a980003753456258244538554b4b4b: UUID="J2hkjU-0X5t-Ntgk-SkHb-3BW2-CqD9-0oRLUb" TYPE="LVM2_member"
/dev/mapper/360a980003753456258244538554b4b49: UUID="wKjYtR-mqwF-iTWg-sn9o-1v3Y-9Sy3-3lGutA" TYPE="LVM2_member"
/dev/mapper/n_vg-nvol: UUID="3c740d52-b8af-4ae9-b476-9754fac6748f" TYPE="ext4"
```

- Because the DMMP devices do not have a file system and because they do not belong to an LVM, you can conclude that they are RAW devices.

1.13.2 Before entering cutover phase

- For FC configurations, ensure that fabric connectivity and zoning to clustered Data ONTAP controllers is established. For iSCSI configurations, ensure that iSCSI sessions to clustered Data ONTAP controllers are discovered and logged in.

2. Stop I/O to the mount points.
3. You must ensure the applications that are accessing the LUNs are shut down gracefully.
4. Flush the DMMP aliases for the 7-Mode LUNs by using the “multipath -f <device_name>” command. See the pretransition steps for identifying the DMMP devices or aliases for the respective 7-Mode LUNs. You can also use the “OS Device ID” column of the Inventory Assessment Workbook’s “SAN Host LUNs” tab to get the DMMP device names.

1.13.3 Posttransition

Posttransition steps begin on the host after the LUNs are transitioned from 7-Mode to clustered Data ONTAP. After the transition, the 7-Mode volumes are in the “offline” state, and the 7-Mode LUNs are not accessible to the host.

After LUN transition, the DMMP device handle ID (WWID) changes; therefore, you must either manually replace the 7-Mode LUN’s WWID with the clustered Data ONTAP LUN’s WWID on all affected files and applications, or you must create a multipath alias with the 7-Mode LUN’s WWID. Both methods are described later in this section.

NOTE: After the transition completes, the 7MTT Tool provides a LUN mapping table between 7-Mode LUNs and their corresponding clustered Data ONTAP LUNs. You can use this table to identify the clustered Data ONTAP LUN device handle IDs and LUN path names.

1. Remove the SCSI devices that were created for the 7-Mode LUNs. See the pretransition steps for identifying the SCSI device IDs. You can use the “rescan-scsi-bus.sh -r” command to remove all the SCSI devices on the host at once, or you can individually delete each SCSI device by using the “echo 1 > /sys/block/<scsi_id>/device/delete” command on all 7-Mode LUN SCSI devices. Use the “SCSI Device ID” column in the Inventory Assessment Workbook’s “SAN Host LUNs” tab to identify the LUN SCSI device IDs.
2. Use the “rescan-scsi-bus.sh” command to discover new clustered Data ONTAP LUNs.
3. Use the “sanlun lun show” command to identify the SCSI devices of the new clustered Data ONTAP LUNs.

Example:

In the example, the SCSI device name for “lun_dmmp_raw” LUN is “/dev/sdz”, and the SCSI device name for “lun_dmmp_raw_with_alias” LUN is “/dev/sdy”:

```
[root@IBMX3550M3-229-167 ~]# sanlun lun show
```

controller(7mode)/ vserver(Cmode)	lun-pathname	device filename	host adapter	protocol	lun size	mode
vs_7MTT	/vol/lun_dmmp_raw_vol/lun_dmmp_raw	/dev/sdz	host10	iSCSI	2g	C
vs_7MTT	/vol/lun_dmmp_raw_with_alias_vol/lun_dmmp_raw_with_alias	/dev/sdy	host10	iSCSI	2g	C
vs_7MTT	/vol/lun_dmmp_raw_vol/lun_dmmp_raw	/dev/sdx	host9	iSCSI	2g	C
vs_7MTT	/vol/lun_dmmp_raw_with_alias_vol/lun_dmmp_raw_with_alias	/dev/sdw	host9	iSCSI	2g	C
vs_7MTT	/vol/lun_dmmp_raw_vol/lun_dmmp_raw	/dev/sdv	host8	iSCSI	2g	C
vs_7MTT	/vol/lun_dmmp_raw_with_alias_vol/lun_dmmp_raw_with_alias	/dev/sdu	host8	iSCSI	2g	C
vs_7MTT	/vol/lun_dmmp_raw_vol/lun_dmmp_raw	/dev/sdh	host7	iSCSI	2g	C
vs_7MTT	/vol/lun_dmmp_raw_with_alias_vol/lun_dmmp_raw_with_alias	/dev/sdq	host7	iSCSI	2g	C
vs_7MTT	/vol/cdot_lun_lvm1_vol/lun_lvm1	/dev/sdt	host10	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm3_vol/lun_lvm3	/dev/sds	host10	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm2_vol/lun_lvm2	/dev/sdr	host10	iSCSI	2g	C
vs_7MTT	/vol/cdot_lun_lvm1_vol/lun_lvm1	/dev/sdq	host9	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm3_vol/lun_lvm3	/dev/sdp	host9	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm2_vol/lun_lvm2	/dev/sdo	host9	iSCSI	2g	C
vs_7MTT	/vol/cdot_lun_lvm1_vol/lun_lvm1	/dev/sdn	host8	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm3_vol/lun_lvm3	/dev/sdm	host8	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm2_vol/lun_lvm2	/dev/sdl	host8	iSCSI	2g	C
vs_7MTT	/vol/cdot_lun_lvm1_vol/lun_lvm1	/dev/sdf	host7	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm3_vol/lun_lvm3	/dev/sde	host7	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm2_vol/lun_lvm2	/dev/sdd	host7	iSCSI	2g	C
vs_7MTT	/vol/cdot_lun_dmmp_fs_without_alias_vol/lun_dmmp_fs_without_alias	/dev/sdk	host10	iSCSI	2g	C
vs_7MTT	/vol/cdot_lun_dmmp_fs_without_alias_vol/lun_dmmp_fs_without_alias	/dev/sdj	host9	iSCSI	2g	C
vs_7MTT	/vol/cdot_lun_dmmp_fs_without_alias_vol/lun_dmmp_fs_without_alias	/dev/sdb	host7	iSCSI	2g	C
vs_7MTT	/vol/cdot_lun_dmmp_fs_without_alias_vol/lun_dmmp_fs_without_alias	/dev/sdi	host8	iSCSI	2g	C
fas3270-229-146	/vol/lun_dmmp_fs_with_alias_vol/lun_dmmp_fs_with_alias	/dev/sdc	host5	iSCSI	2g	7

4. Obtain the “wwid” for the new clustered Data ONTAP LUNs by running “/lib/udev/scsi_id -g -u -d /dev/<SCSI dev>” command. Or, to locate the new clustered Data ONTAP device handle ID, see the “LUN mapping table” that is generated by the 7MTT tool.

Example:

In the example, the “wwid” for SCSI “/dev/sdz”, which belongs to LUN “lun_dmmp_raw”, is “3600a098051764937303f4479515a744e”, and the “wwid” for SCSI “/dev/sdy”, which belongs to LUN “lun_dmmp_raw_with_alias” is “3600a098051764937303f4479515a744d”:

```
[root@IBMX3550M3-229-167 ~]# /lib/udev/scsi_id -g -u -d /dev/sdz
3600a098051764937303f4479515a744e
[root@IBMX3550M3-229-167 ~]# /lib/udev/scsi_id -g -u -d /dev/sdy
3600a098051764937303f4479515a744d
```

5. Update the “/etc/multipath.conf” file to replace the 7-Mode device handler ID with its corresponding clustered Data ONTAP device handle ID so that the DMMP alias points to clustered Data ONTAP LUN.

Example:

The example depicts two DMMP RAW devices; one with the alias “raw_device”, and the other without an alias, but rather a device handler ID. The “/etc/multipath.conf” file is updated for both scenarios.

In the first scenario, the DMMP alias “raw_device” is modified to point to clustered Data ONTAP device handle ID.

In the second scenario, a new DMMP alias is created with the 7-Mode LUN’s device handle ID to reference the clustered Data ONTAP LUN’s device handle ID.

This output displays the /etc/multipath.conf file before the update:

```
[root@IBMX3550M3-229-167 ~]# cat /etc/multipath.conf
blacklist {
    wwid 3600605b00425863016fe22f30862c5b1
}

multipaths {
    multipath {
        wwid 360a980003753456258244538554b4b4f
        alias raw_device
    }
}
```

This output displays the /etc/multipath.conf file after the update:

```

[root@IBMX3550M3-229-167 ~]# cat /etc/multipath.conf
blacklist {
    wwid 3600605b00425863016fe22f30862c5b1
}

multipaths {
    multipath {
        wwid 360a980003753456258244538554b4b4f
        alias raw_device
    }
}

[root@IBMX3550M3-229-167 ~]# vim /etc/multipath.conf
[root@IBMX3550M3-229-167 ~]# vim /etc/multipath.conf
[root@IBMX3550M3-229-167 ~]# cat /etc/multipath.conf
blacklist {
    wwid 3600605b00425863016fe22f30862c5b1
}

multipaths {

    # RAW device with alias, updated WWID of new cDOT LUN
    multipath {
        wwid 3600a098051764937303f4479515a744d
        alias raw_device
    }

    # Created new alias with old WWID LUN name for new cDOT LUN
    multipath {
        wwid 3600a098051764937303f4479515a744e
        alias 360a980003753456258244538554b4b4d
    }
}

```

6. Run the “multipath” command to configure DMMP devices.

1.14 DMMP device with file system

1. From the “LUN Name” column of the Inventory Assessment Workbook’s “LUN” tab, identify the name of the LUN to be transitioned.
2. From the “SCSI Device ID” column of the Inventory Assessment Workbook’s “SAN Host LUNs” tab, locate the SCSI device names for the LUN to be transitioned.

Alternatively, use the “sanlun lun show” command to identify the SCSI device name for a LUN.

Example:

In this example, the transition LUN is “lun_dmmp_fs_without_alias” and the respective SCSI device is “/dev/sdb”:

```
[root@IBMX3550M3-229-167 mapper]# sanlun lun show
```

controller (7mode)/ vserver (Cmode)	lun-pathname	device filename	host adapter	protocol	lun size	mode
fas3270-229-147	/vol/lun_dmmp_raw_with_alias_vol/lun_dmmp_raw_with_alias	/dev/sdh	host6	iSCSI	2g	7
fas3270-229-147	/vol/lun_dmmp_raw_vol/lun_dmmp_raw	/dev/sdg	host6	iSCSI	2g	7
fas3270-229-147	/vol/lun_lvm3_vol/lun_lvm3	/dev/sdf	host6	iSCSI	2g	7
fas3270-229-147	/vol/lun_lvm2_vol/lun_lvm2	/dev/sde	host6	iSCSI	2g	7
fas3270-229-146	/vol/lun_lvm1_vol/lun_lvm1	/dev/sdd	host5	iSCSI	2g	7
fas3270-229-146	/vol/lun_dmmp_fs_with_alias_vol/lun_dmmp_fs_with_alias	/dev/sdc	host5	iSCSI	2g	7
fas3270-229-146	/vol/lun_dmmp_fs_without_alias_vol/lun_dmmp_fs_without_alias	/dev/sdb	host5	iSCSI	2g	7

- From the “OS Device ID”.column of the Inventory Assessment Workbook’s “SAN Host LUNs” tab, identify the DMMP device name for the transition LUN.

Alternatively, you can use the “multipath -ll” command to identify the DMMP device for the LUN. The DMMP device name is either a unique device handle ID (WWID) or an alias.

Example:

This example indicates that the DMMP device name for SCSI device “/dev/sdb” is “360a9800037534562572b453855496b31”:

```
[root@IBMX3550M3-229-167 mapper]# multipath -ll
360a9800037534562572b453855496b31 dm-3 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
+- policy='round-robin 0' prio=2 status=active
  - 5:0:0:0 sdb 8:16 active ready running
```

- Verify that the DMMP device exists in the “/dev/mapper” directory by running the “ls /dev/mapper/ <DMMP device name>” command. If you cannot locate the device, the DMMP device might be using an alias or a user-friendly name.

Example:

The output for the “ls /dev/mapper/360a9800037534562572b453855496b31” indicates that the DMMP device exists:

```
[root@IBMX3550M3-229-167 ~]# ls /dev/mapper/360a9800037534562572b453855496b31
/dev/mapper/360a9800037534562572b453855496b31
```

- From the Inventory Assessment Workbook’s “SAN Host LVMs” tab, ensure that the LUN identified for transition does not have LVM configured. If the LV device entry is not present then LVM is not configured on that device.

Alternatively, you can determine whether the DMMP device is part of an LVM by running “blkid” command. Look at the TYPE value of the DMMP device. If the TYPE is “LVM2_member”, the DMMP is part of an LVM.

Example:

In the sample output, DMMP device “360a9800037534562572b453855496b31” is not part of an LVM:

```
[root@IBMX3550M3-229-167 mapper]# blkid
/dev/mapper/vg_ibmx3550m3229167-lv_root: UUID="4119733f-b52e-4bef-bbcb-2d2a623b6fdc" TYPE="ext4"
/dev/sda1: UUID="c875432c-1f86-4ac0-9cb7-5d832b23c377" TYPE="ext4"
/dev/sda2: UUID="hSq1V6-JrvF-Ppfi-kmYM-Haf0-RAx6-Q2I8Ho" TYPE="LVM2_member"
/dev/mapper/vg_ibmx3550m3229167-lv_swap: UUID="77014c28-2e68-4881-8b11-f24d4de58440" TYPE="swap"
/dev/mapper/vg_ibmx3550m3229167-lv_home: UUID="119e8728-d544-4572-a3f0-5c0c0ecf3e7f" TYPE="ext4"
/dev/mapper/360a9800037534562572b453855496b31: UUID="a21b6900-c8f8-4d5c-b97c-9528c999b3e7" TYPE="ext4"
/dev/mapper/360a9800037534562572b453855496b35: UUID="dLp0zP-PYCo-LiY1-W346-BZi8-zL5K-gQzKHd" TYPE="LVM2_member"
/dev/mapper/360a980003753456258244538554b4b4b: UUID="J2hkjU-0X5t-Ntgk-SkHb-3BW2-CqD9-0oRLUb" TYPE="LVM2_member"
/dev/mapper/360a980003753456258244538554b4b49: UUID="wKjYtR-mqwF-iTWg-sn9o-1v3Y-9Sy3-3lGutA" TYPE="LVM2_member"
/dev/mapper/n_vg-nvol: UUID="3c740d52-b8af-4ae9-b476-9754fac6748f" TYPE="ext4"
```

- From the Inventory Assessment Workbook’s “SAN Host File System” tab, identify the file system configured on the DMMP device.

Alternatively, you can use running the “blkid” command to determine whether the DMMP device has a file system. Check whether the TYPE value of the DMMP device indicates a file system exists.

Example:

In the sample output, DMMP device “360a9800037534562572b453855496b31” indicates the TYPE value is “ext4”, which is a file system:

```
[root@IBMX3550M3-229-167 mapper]# blkid
/dev/mapper/vg_ibmx3550m3229167-lv_root: UUID="4119733f-b52e-4bef-bbcb-2d2a623b6fdc" TYPE="ext4"
/dev/sda1: UUID="c875432c-1f86-4ac0-9cb7-5d832b23c377" TYPE="ext4"
/dev/sda2: UUID="hSq1V6-JrvF-Ppfi-kmYM-Haf0-RAX6-Q2I8Ho" TYPE="LVM2_member"
/dev/mapper/vg_ibmx3550m3229167-lv_swap: UUID="77014c28-2e68-4881-8b11-f24d4de58440" TYPE="swap"
/dev/mapper/vg_ibmx3550m3229167-lv_home: UUID="119e8728-d544-4572-a3f0-5c0c0ecf3e7f" TYPE="ext4"
/dev/mapper/360a9800037534562572b453855496b31: UUID="a21b6900-c8f8-4d5c-b97c-9528c999b3e7" TYPE="ext4"
/dev/mapper/360a9800037534562572b453855496b35: UUID="dLp0zP-PYCo-LiY1-W346-BZi8-zL5K-gQzKHd" TYPE="LVM2_member"
/dev/mapper/360a980003753456258244538554b4b4b: UUID="J2hkjU-0X5t-Ntgk-SkHb-3BW2-CqD9-0oRLUb" TYPE="LVM2_member"
/dev/mapper/360a980003753456258244538554b4b49: UUID="wKjYtR-mqWf-iTWg-sn9o-1v3Y-9Sy3-3lGutA" TYPE="LVM2_member"
/dev/mapper/n_vg-nvol: UUID="3c740d52-b8af-4ae9-b476-9754fac6748f" TYPE="ext4"
```

7. From the “UUID” column of the Inventory Assessment Workbook’s “SAN Host File System” tab, identify the file system UUID of the LUN.

Alternatively, use the “blkid” command to identify the file system “UUID” of the LUN.

Example:

In the sample output, DMMP device “360a9800037534562572b453855496b31” has an “ext4” file system, and its UUID is “a21b6900-c8f8-4d5c-b97c-9528c999b3e7”:

```
[root@IBMX3550M3-229-167 mapper]# blkid
/dev/mapper/vg_ibmx3550m3229167-lv_root: UUID="4119733f-b52e-4bef-bbcb-2d2a623b6fdc" TYPE="ext4"
/dev/sda1: UUID="c875432c-1f86-4ac0-9cb7-5d832b23c377" TYPE="ext4"
/dev/sda2: UUID="hSq1V6-JrvF-Ppfi-kmYM-Haf0-RAX6-Q2I8Ho" TYPE="LVM2_member"
/dev/mapper/vg_ibmx3550m3229167-lv_swap: UUID="77014c28-2e68-4881-8b11-f24d4de58440" TYPE="swap"
/dev/mapper/vg_ibmx3550m3229167-lv_home: UUID="119e8728-d544-4572-a3f0-5c0c0ecf3e7f" TYPE="ext4"
/dev/mapper/360a9800037534562572b453855496b31: UUID="a21b6900-c8f8-4d5c-b97c-9528c999b3e7" TYPE="ext4"
/dev/mapper/360a9800037534562572b453855496b35: UUID="dLp0zP-PYCo-LiY1-W346-BZi8-zL5K-gQzKHd" TYPE="LVM2_member"
/dev/mapper/360a980003753456258244538554b4b4b: UUID="J2hkjU-0X5t-Ntgk-SkHb-3BW2-CqD9-0oRLUb" TYPE="LVM2_member"
/dev/mapper/360a980003753456258244538554b4b49: UUID="wKjYtR-mqWf-iTWg-sn9o-1v3Y-9Sy3-3lGutA" TYPE="LVM2_member"
/dev/mapper/n_vg-nvol: UUID="3c740d52-b8af-4ae9-b476-9754fac6748f" TYPE="ext4"
```

8. From the “Mount” column of the “SAN Host Filesystems” tab, identify the directory on which the DMMP device is mounted.

Alternatively, use either the “df -h” command of the “mount” command to determine whether the DMMP device is mounted.

Example:

The sample output indicates that DMMP device “360a9800037534562572b453855496b31” is mounted on “/mnt/dmmp_fs”:

```
[root@IBMX3550M3-229-167 mapper]# mount
/dev/mapper/vg_ibmx3550m3229167-lv_root on / type ext4 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
tmpfs on /dev/shm type tmpfs (rw,rootcontext="system_u:object_r:tmpfs_t:s0")
/dev/sda1 on /boot type ext4 (rw)
/dev/mapper/vg_ibmx3550m3229167-lv_home on /home type ext4 (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
/dev/mapper/n_vg-nvol on /mnt/LVM_Mount type ext4 (rw,_netdev)
/dev/mapper/360a9800037534562572b453855496b31 on /mnt/dmmp_fs type ext4 (rw,_netdev)
```

9. Determine whether a mount point entry for the DMMP device exists in the “/etc/fstab” file.

Example:

The sample “/etc/fstab” output indicates that a mount point entry exists for DMMP device “360a9800037534562572b453855496b31”:

```
[root@IBMX3550M3-229-167 mapper]# cat /etc/fstab
/dev/mapper/vg_ibmx3550m3229167-lv_root / ext4 defaults 1 1
UUID=c875432c-1f86-4ac0-9cb7-5d832b23c377 /boot ext4 defaults 1 2
/dev/mapper/vg_ibmx3550m3229167-lv_home /home ext4 defaults 1 2
/dev/mapper/vg_ibmx3550m3229167-lv_swap swap swap defaults 0 0
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/n_vg-nvol /mnt/LVM_Mount ext4 defaults,_netdev 0 0
/dev/mapper/360a9800037534562572b453855496b31 /mnt/dmmp_fs ext4 defaults,_netdev 0 0
```

10. Make a backup of the “/etc/fstab” file before editing it.

Example:

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

11. Edit the “/etc/fstab” file and replace the DMMP device handle ID with the file system UUID.

Example:

In this example, the “/etc/fstab” file is edited to replace DMMP device “/dev/mapper/360a9800037534562572b453855496b31” with “UUID=a21b6900-c8f8-4d5c-b97c-9528c999b3e7”:

```
[root@IBMX3550M3-229-167 mapper]# cat /etc/fstab
/dev/mapper/vg_ibmx3550m3229167-lv_root / ext4 defaults 1 1
UUID=c875432c-1f86-4ac0-9cb7-5d832b23c377 /boot ext4 defaults 1 2
/dev/mapper/vg_ibmx3550m3229167-lv_home /home ext4 defaults 1 2
/dev/mapper/vg_ibmx3550m3229167-lv_swap swap swap defaults 0 0
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/n_vg-nvol /mnt/LVM_Mount ext4 defaults,_netdev 0 0
UUID=a21b6900-c8f8-4d5c-b97c-9528c999b3e7 /mnt/dmmp_fs ext4 defaults,_netdev 0 0
```

1.14.1 Before entering cutover phase

1. For FC configurations, ensure that fabric connectivity and zoning to clustered Data ONTAP controllers is established. For iSCSI configurations, ensure that iSCSI sessions to clustered Data ONTAP controllers are discovered and logged in.
2. Stop IO on all the mount points.
3. Ensure that the applications that are accessing the LUNs are gracefully shut down.
4. Unmount all mount points.
5. Flush the DMMP aliases for 7-Mode LUNs using the “multipath -f <device_name>” command. See the pretransition steps for identifying the DMMP devices or aliases for the respective 7-Mode LUNs. Or, get the DMMP device names from the “OS Device ID” column in the Inventory Assessment Workbook’s “SAN Host LUNs” tab.

1.14.2 Posttransition

Posttransition steps begin on the host after the LUNs are transitioned to clustered Data ONTAP. After the transition, the 7-Mode volumes are in the “offline” state, and the 7-Mode LUNs are not accessible to the host.

After LUN transition the DMMP device handle ID (WWID) changes.

NOTE: After the transition completes, the 7MTT Tool provides a LUN mapping table between 7-Mode LUNs and the corresponding clustered Data ONTAP LUNs. You can use this table to get clustered Data ONTAP LUN device handles ID and LUN path names.

1. Remove the SCSI devices that were created for the 7-Mode LUNs. See the pretransition steps for identifying the SCSI device IDs. You can use the “rescan-scsi-bus.sh -r” command to remove all the SCSI devices on the host at once, or you can individually delete each SCSI device by using the “echo 1 > /sys/block/<scsi_id>/device/delete” command on all 7-Mode LUN SCSI devices. Use the “SCSI Device ID” column in the Inventory Assessment Workbook’s “SAN Host LUNs” tab to identify the LUN SCSI device IDs.
2. Run the “rescan-scsi-bus.sh” command to discover new clustered Data ONTAP LUNs.
3. Run the “multipath” command to apply multipath configurations to newly discovered clustered Data ONTAP LUNs.
4. Run the “sanlun lun show” command to verify the clustered Data ONTAP LUNs are present:

```
[root@IBMx3550M3-229-167 ~]# sanlun lun show
```

controller (7mode)/ vserver (Cmode)	lun-pathname	device filename	host adapter	protocol	lun size	mode
vs_7MITT	/vol/cdot_lun_dmmp_fs_without_alias_vol/lun_dmmp_fs_without_alias	/dev/sdk	host10	iSCSI	2g	C
vs_7MITT	/vol/cdot_lun_dmmp_fs_without_alias_vol/lun_dmmp_fs_without_alias	/dev/sdj	host9	iSCSI	2g	C
vs_7MITT	/vol/cdot_lun_dmmp_fs_without_alias_vol/lun_dmmp_fs_without_alias	/dev/sdb	host7	iSCSI	2g	C
vs_7MITT	/vol/cdot_lun_dmmp_fs_without_alias_vol/lun_dmmp_fs_without_alias	/dev/sdi	host8	iSCSI	2g	C
fas3270-229-147	/vol/lun_dmmp_raw_with_alias_vol/lun_dmmp_raw_with_alias	/dev/sdh	host6	iSCSI	2g	7
fas3270-229-147	/vol/lun_dmmp_raw_vol/lun_dmmp_raw	/dev/sdg	host6	iSCSI	2g	7
fas3270-229-147	/vol/lun_lvm3_vol/lun_lvm3	/dev/sdf	host6	iSCSI	2g	7
fas3270-229-147	/vol/lun_lvm2_vol/lun_lvm2	/dev/sde	host6	iSCSI	2g	7
fas3270-229-146	/vol/lun_lvm1_vol/lun_lvm1	/dev/sdd	host5	iSCSI	2g	7
fas3270-229-146	/vol/lun_dmmp_fs_with_alias_vol/lun_dmmp_fs_with_alias	/dev/sdc	host5	iSCSI	2g	7

5. Run the “multipath -ll” command to verify DMMP devices are present:

```
[root@IBMX3550M3-229-167 ~]# multipath -ll
360a980003753456258244538554b4bf dm-8 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
+- policy='round-robin 0' prio=2 status=active
  - 6:0:0:3 sdh 8:112 active ready running
360a9800037534562572b453855496b35 dm-5 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
+- policy='round-robin 0' prio=2 status=active
  - 5:0:0:2 sdd 8:48 active ready running
360a980003753456258244538554b4bd dm-4 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
+- policy='round-robin 0' prio=2 status=active
  - 6:0:0:2 sdg 8:96 active ready running
360a9800037534562572b453855496b33 dm-3 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
+- policy='round-robin 0' prio=2 status=active
  - 5:0:0:1 sdc 8:32 active ready running
360a980003753456258244538554b4bb dm-6 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
+- policy='round-robin 0' prio=2 status=active
  - 6:0:0:1 sdf 8:80 active ready running
360a980003753456258244538554b4b49 dm-7 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
+- policy='round-robin 0' prio=2 status=active
  - 6:0:0:0 sde 8:64 active ready running
3600a098051764937303f4479515a744b dm-9 NETAPP,LUN C-Mode
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='1 alua' wp=rw
+- policy='round-robin 0' prio=50 status=active
  | 8:0:0:0 sdi 8:128 active ready running
  | 10:0:0:0 sdk 8:160 active ready running
+- policy='round-robin 0' prio=10 status=enabled
  | 7:0:0:0 sdb 8:16 active ready running
  | 9:0:0:0 sdj 8:144 active ready running
```

6. If the mount points are defined in the “/etc/fstab” file, run either the “mount -a” command to mount all the mount points, or run the “mount <device_name> <mountpoint>” command to manually mount each device.
7. Verify the mount points by running the “mount” command:

```
[root@IBMX3550M3-229-167 ~]# mount
/dev/mapper/vg_ibmx3550m3229167-lv_root on / type ext4 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
tmpfs on /dev/shm type tmpfs (rw,rootcontext="system_u:object_r:tmpfs_t:s0")
/dev/sdal on /boot type ext4 (rw)
/dev/mapper/vg_ibmx3550m3229167-lv_home on /home type ext4 (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
/dev/mapper/n_vg-nvol on /mnt/LVM_Mount type ext4 (rw,_netdev)
/dev/mapper/3600a098051764937303f4479515a744b on /mnt/dmmp_fs type ext4 (rw,_netdev)
```

1.15 File system on LVM device

1.15.1 Pretransition

1. From the “LUN Name” column of the Inventory Assessment Workbook’s “LUN” tab, identify the name of the LUN being transitioned.
2. From the “SCSI Device ID” column of the Inventory Assessment Workbook’s “SAN Host LUNs” tab, identify the SCSI device names for the LUNs being transitioned.

Alternatively, use the “sanlun lun show” command to determine the SCSI device name for a LUN.

Example:

In this example, the transition LUNs are “lun_lvm1”, “lun_lvm2” and “lun_lvm3” and their respective SCSI devices are “/dev/sdd”, “/dev/sde” and “/dev/sdf”:

```
[root@IBMX3550M3-229-167 ~]# sanlun lun show
```

controller(7mode)/ vserver(Cmode)	lun-pathname	device filename	host adapter	protocol	lun size	mode
fas3270-229-147	/vol/lun_dmmp_raw_with_alias_vol/lun_dmmp_raw_with_alias	/dev/sdh	host6	iSCSI	2g	7
fas3270-229-147	/vol/lun_dmmp_raw_vol/lun_dmmp_raw	/dev/sdg	host6	iSCSI	2g	7
fas3270-229-147	/vol/lun_lvm3_vol/lun_lvm3	/dev/sdf	host6	iSCSI	2g	7
fas3270-229-147	/vol/lun_lvm2_vol/lun_lvm2	/dev/sde	host6	iSCSI	2g	7
fas3270-229-146	/vol/lun_lvm1_vol/lun_lvm1	/dev/sdd	host5	iSCSI	2g	7
fas3270-229-146	/vol/lun_dmmp_fs_with_alias_vol/lun_dmmp_fs_with_alias	/dev/sdc	host5	iSCSI	2g	7
fas3270-229-146	/vol/lun_dmmp_fs_without_alias_vol/lun_dmmp_fs_without_alias	/dev/sdb	host5	iSCSI	2g	7

- From the “OS Device ID” column of the Inventory Assessment Workbook’s “SAN Host LUNs” tab, identify the DMMP device name for the LUN being transitioned.

Alternatively, use the the “multipath -ll” command to determine the DMMP device name for the LUN. The DMMP device name can be either a unique device handle ID (WWID) or an alias.

Example:

In this example, the DMMP device name for SCSI device “/dev/sdd” is “360a9800037534562572b453855496b35”, the DMMP device name for SCSI device “/dev/sde” is “360a980003753456258244538554b4b49”, and the DMMP device name for SCSI device “/dev/sdf” is “360a980003753456258244538554b4b4b”:

```
[root@IBMX3550M3-229-167 ~]# multipath -ll
360a980003753456258244538554b4b4f dm-9 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
'++ policy='round-robin 0' prio=2 status=active
  '- 6:0:0:3 sdh 8:112 active ready running
360a9800037534562572b453855496b35 dm-5 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
'++ policy='round-robin 0' prio=2 status=active
  '- 5:0:0:2 sdd 8:48 active ready running
360a980003753456258244538554b4b4d dm-8 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
'++ policy='round-robin 0' prio=2 status=active
  '- 6:0:0:2 sdg 8:96 active ready running
360a9800037534562572b453855496b33 dm-4 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
'++ policy='round-robin 0' prio=2 status=active
  '- 5:0:0:1 sdc 8:32 active ready running
360a980003753456258244538554b4b4b dm-6 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
'++ policy='round-robin 0' prio=2 status=active
  '- 6:0:0:1 sdf 8:80 active ready running
360a9800037534562572b453855496b31 dm-3 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
'++ policy='round-robin 0' prio=2 status=active
  '- 5:0:0:0 sdb 8:16 active ready running
360a980003753456258244538554b4b49 dm-7 NETAPP,LUN
size=2.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
'++ policy='round-robin 0' prio=2 status=active
  '- 6:0:0:0 sde 8:64 active ready running
```

- Verify that the DMMP device in the “/dev/mapper” directory exists by running the “ls /dev/mapper/ <DMMP device name>” command. If you cannot locate the device, the DMMP device might be using an alias or user-friendly name.

Example:

In this example, “ls /dev/mapper/360a9800037534562572b453855496b35” indicates that the DMMP device exists in the /dev/mapper directory:

```
[root@IBMX3550M3-229-167 ~]# ls /dev/mapper/360a9800037534562572b453855496b35
/dev/mapper/360a9800037534562572b453855496b35
[root@IBMX3550M3-229-167 ~]# ls /dev/mapper/360a980003753456258244538554b4b49
/dev/mapper/360a980003753456258244538554b4b49
[root@IBMX3550M3-229-167 ~]# ls /dev/mapper/360a980003753456258244538554b4b4b
/dev/mapper/360a980003753456258244538554b4b4b
```

- From the Inventory Assessment Workbook's "SAN Host LVMs" tab, ensure that the LUN identified for transition has LVM configured. If the LV device entry exists, then LVM is configured on the device.

Alternatively, use the running "blkid" command to determine whether the DMMP device is part of an LVM. The TYPE value displayed for the DMMP device indicates whether it is an LVM member. For example, if the TYPE is "LVM2_member" then the DMMP device is part of an LVM.

Example:

In this example, DMMP device "360a9800037534562572b453855496b35", "360a980003753456258244538554b4b4b", and "360a980003753456258244538554b4b49" are part of an LVM:

```
[root@IBMX3550M3-229-167 mapper]# blkid
/dev/mapper/vg_ibmx3550m3229167-lv_root: UUID="4119733f-b52e-4bef-bbcb-2d2a623b6fdc" TYPE="ext4"
/dev/sda1: UUID="c875432c-1f86-4ac0-9cb7-5d832b23c377" TYPE="ext4"
/dev/sda2: UUID="hSq1V6-JrvF-Ppfi-kmYM-Haf0-RAX6-Q2I8Ho" TYPE="LVM2_member"
/dev/mapper/vg_ibmx3550m3229167-lv_swap: UUID="77014c28-2e68-4881-8b11-f24d4de58440" TYPE="swap"
/dev/mapper/vg_ibmx3550m3229167-lv_home: UUID="119e8728-d544-4572-a3f0-5c0c0ecf3e7f" TYPE="ext4"
/dev/mapper/360a9800037534562572b453855496b31: UUID="a21b6900-c8f8-4d5c-b97c-9528c999b3e7" TYPE="ext4"
/dev/mapper/360a9800037534562572b453855496b35: UUID="dLp0zP-PYCo-LiY1-W346-BZi8-zL5K-gQzKHd" TYPE="LVM2_member"
/dev/mapper/360a980003753456258244538554b4b4b: UUID="J2hkjU-0X5t-Ntgk-SkHb-3BW2-CqD9-0cRLUb" TYPE="LVM2_member"
/dev/mapper/360a980003753456258244538554b4b49: UUID="wKjYtR-mqwF-iTWg-sn9c-1v3Y-9Sy3-3lGutA" TYPE="LVM2_member"
/dev/mapper/n_vg-nvol: UUID="3c740d52-b8af-4ae9-b476-9754fac6748f" TYPE="ext4"
```

- From the "Physical Volume Name", the "Volume Group Name", and the "Logical Volume Path" columns, of the Inventory Assessment Workbook's "SAN Host LVMs" tab, identify the Physical Volume (PV), Volume Group (VG), and the Logical Volume (LV), respectively.
- No pretransition steps are required if the DMMP device is part of an LVM.

1.15.2 Before entering cutover phase

- If you are using an FC configuration, you must ensure that fabric connectivity and zoning to clustered Data ONTAP controllers is established. If you are using an iSCSI configuration, you must ensure that iSCSI sessions to clustered Data ONTAP controllers are discovered and logged in.
- Stop IO on all the mount points.
- Ensure that the applications that are accessing the LUNs are gracefully shut down.
- Unmount all the mount points.
- Run the "vgchange -an <vg_name>" command to disable volume groups:

```
[root@IBMX3550M3-229-167 ~]# vgchange -an n_vg
0 logical volume(s) in volume group "n_vg" now active
```

- Verify the Logical Volume (LV) status using the "lvdisplay" command. The LV status displays as "NOT available".

Example:

```
[root@IBMx3550M3-229-167 ~]# vgchange -an n_vg
0 logical volume(s) in volume group "n_vg" now active
[root@IBMx3550M3-229-167 ~]#
[root@IBMx3550M3-229-167 ~]# lvdisplay
--- Logical volume ---
LV Path                /dev/n_vg/nvol
LV Name                 nvol
VG Name                n_vg
LV UUID                IrNKYy-lo3y-Ds50-2I51-vEVj-7X7K-Et9Lnu
LV Write Access         read/write
LV Creation host, time IBMx3550M3-229-167.gdl.englab.netapp.com, 2014-07-30 04:40:27 -0400
LV Status               NOT available
LV Size                3.00 GiB
Current LE              768
Segments                1
Allocation              inherit
Read ahead sectors     auto
```

7. Run the “vgexport” command to export the VG:

```
[root@IBMx3550M3-229-167 ~]# vgexport n_vg
Volume group "n_vg" successfully exported
```

8. Verify the VG export status by running the “vgdisplay” command. The VG status displays as “exported”:

```
[root@IBMx3550M3-229-167 ~]# vgdisplay
Volume group n_vg is exported
--- Volume group ---
VG Name                n_vg
System ID
Format                 lvm2
Metadata Areas          3
Metadata Sequence No    3
VG Access               read/write
VG Status               exported/resizable
MAX LV                  0
Cur LV                 1
Open LV                 0
Max PV                  0
Cur PV                 3
Act PV                  3
VG Size                 5.99 GiB
PE Size                 4.00 MiB
Total PE                1533
Alloc PE / Size         768 / 3.00 GiB
Free PE / Size          765 / 2.99 GiB
VG UUID                 5JL0LH-tHgV-frxo-pgVE-peb2-iS1a-lubncl
```

9. Flush the DMMP names for the 7-Mode LUNs using the “multipath -f <device_name>” command. See the pretransition steps for identifying the DMMP devices or aliases for the respective 7-Mode LUNs. Or, use the “OS Device ID” column of the Inventory Assessment Workbook’s “SAN Host LUNs” tab to identify the DMMP device names.

1.15.3 Posttransition

Posttransition steps begin on the host after the LUNs are transitioned to clustered Data ONTAP.

After the transition, the 7-Mode volumes are in the “offline” state, and the 7-Mode LUNs are not accessible to the host.

After LUN transition, the LVM attributes, for example, the LV and VG names, do not change, so you will continue to use those LV and VG names for posttransition configuration.

NOTE: After the transition completes, the 7MTT Tool provides a LUN mapping table between 7-Mode LUNs and their corresponding clustered Data ONTAP LUNs. You can use this table to identify the clustered Data ONTAP LUN device handle IDs and LUN path names.

1. Remove the SCSI devices that were created for the 7-Mode LUNs. See the pretransition steps for identifying the SCSI device IDs. You can use the “rescan-scsi-bus.sh -r” command to remove all the SCSI devices on the host at once, or you can individually delete each SCSI device by using the “echo 1 > /sys/block/<scsi_id>/device/delete” command on all 7-Mode LUN SCSI devices. Use the “SCSI Device ID” column in the Inventory Assessment Workbook’s “SAN Host LUNs” tab to identify the LUN SCSI device IDs.
2. Run the “rescan-scsi-bus.sh” command to discover new clustered Data ONTAP LUNs.
3. Run the “multipath” command to reconfigure DMMP devices.
4. Verify new LUNs by running the “sanlun lun show” command:

```
[root@IBMX3550M3-229-167 ~]# sanlun lun show
```

controller(7mode)/ vserver(Cmode)	lun-pathname	device filename	host adapter	protocol	lun size	mode
vs_7MTT	/vol/cdot_lun_lvm1_vol/lun_lvm1	/dev/sdt	host10	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm3_vol/lun_lvm3	/dev/sds	host10	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm2_vol/lun_lvm2	/dev/sdr	host10	iSCSI	2g	C
vs_7MTT	/vol/cdot_lun_lvm1_vol/lun_lvm1	/dev/sdq	host9	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm3_vol/lun_lvm3	/dev/sdp	host9	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm2_vol/lun_lvm2	/dev/sdo	host9	iSCSI	2g	C
vs_7MTT	/vol/cdot_lun_lvm1_vol/lun_lvm1	/dev/sdn	host8	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm3_vol/lun_lvm3	/dev/sdm	host8	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm2_vol/lun_lvm2	/dev/sdl	host8	iSCSI	2g	C
vs_7MTT	/vol/cdot_lun_lvm1_vol/lun_lvm1	/dev/sdf	host7	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm3_vol/lun_lvm3	/dev/sde	host7	iSCSI	2g	C
vs_7MTT	/vol/lun_lvm2_vol/lun_lvm2	/dev/sdd	host7	iSCSI	2g	C
vs_7MTT	/vol/cdot_lun_dmmp_fs_without_alias_vol/lun_dmmp_fs_without_alias	/dev/sdk	host10	iSCSI	2g	C
vs_7MTT	/vol/cdot_lun_dmmp_fs_without_alias_vol/lun_dmmp_fs_without_alias	/dev/sdj	host9	iSCSI	2g	C
vs_7MTT	/vol/cdot_lun_dmmp_fs_without_alias_vol/lun_dmmp_fs_without_alias	/dev/sdb	host7	iSCSI	2g	C
vs_7MTT	/vol/cdot_lun_dmmp_fs_without_alias_vol/lun_dmmp_fs_without_alias	/dev/sdi	host8	iSCSI	2g	C
fas3270-229-147	/vol/lun_dmmp_raw_with_alias_vol/lun_dmmp_raw_with_alias	/dev/sdh	host6	iSCSI	2g	7
fas3270-229-147	/vol/lun_dmmp_raw_vol/lun_dmmp_raw	/dev/sdg	host6	iSCSI	2g	7
fas3270-229-146	/vol/lun_dmmp_fs_with_alias_vol/lun_dmmp_fs_with_alias	/dev/sdc	host5	iSCSI	2g	7

5. Import the VG by running the “vgimport <vg_name>”:

```
[root@IBMX3550M3-229-167 ~]# vgimport n_vg
Volume group "n_vg" successfully imported
```

6. Verify the VG status by running the “vgdisplay” command:

```
[root@IBMX3550M3-229-167 ~]# vgdisplay
--- Volume group ---
VG Name                n_vg
System ID
Format                 lvm2
Metadata Areas         3
Metadata Sequence No   4
VG Access              read/write
VG Status              resizable
MAX LV                 0
Cur LV                1
Open LV               0
Max PV                 0
Cur PV                3
Act PV                3
VG Size                5.99 GiB
PE Size                4.00 MiB
Total PE              1533
Alloc PE / Size       768 / 3.00 GiB
Free PE / Size        765 / 2.99 GiB
VG UUID                5JL0LH-tHgV-frxo-pgVE-peb2-iSla-lubnc1
```

7. Run the “vgchange -ay <vg_name>” to activate all the LVs:

```
[root@IBMX3550M3-229-167 ~]# vgchange -ay n_vg
1 logical volume(s) in volume group "n_vg" now active
```

8. Verify the LV status by running the “lvdisplay” command:

```
[root@IBMX3550M3-229-167 ~]# lvdisplay
--- Logical volume ---
LV Path                /dev/n_vg/nvol
LV Name                nvol
VG Name                n_vg
LV UUID                IrNKYy-1o3y-Ds50-ZI51-vEVj-7X7K-Et9Lnu
LV Write Access        read/write
LV Creation host, time IBMx3550M3-229-167.gdl.englab.netapp.com, 2014-07-30 04:40:27 -0400
LV Status              available
# open                 0
LV Size                3.00 GiB
Current LE             768
Segments               1
Allocation              inherit
Read ahead sectors     auto
- currently set to     768
Block device           253:10
```

9. If the mount points are defined in the “/etc/fstab” file, run either the “mount -a” command to mount all the mount points, or use the “mount <device_name> <mountpoint>” command to manually mount each of the devices.
10. Run the “mount” command to verify the mount points:

```
[root@IBMX3550M3-229-167 ~]# mount
/dev/mapper/vg_ibmx3550m3229167-lv_root on / type ext4 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
tmpfs on /dev/shm type tmpfs (rw,rootcontext="system_u:object_r:tmpfs_t:s0")
/dev/sda1 on /boot type ext4 (rw)
/dev/mapper/vg_ibmx3550m3229167-lv_home on /home type ext4 (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
/dev/mapper/3600a098051764937303f4479515a744b on /mnt/dmmp_fs type ext4 (rw,_netdev)
/dev/mapper/n_vg-nvol on /mnt/LVM_Mount type ext4 (rw,_netdev)
```

SAN Boot LUN Scenarios

1.16 Protocol: FC/FCoE

1.16.1 Pretransition

1. From the 7-Mode controller's console, identify the 7-Mode LUN name on which RHEL 6 OS is installed. See the "LUN Name" column of the Inventory Assessment Workbook's "LUN" tab. Alternatively, on the host, use the 7-Mode command "lun show" to list the LUNs on controller.
2. Get the SCSI device names for the transition LUN from the Inventory Assessment Workbook's "SAN Host LUNs" tab. Under this tab refer column name "SCSI Device ID". Alternatively, you may get SCSI device name for a given LUN on host by running "sanlun lun show" command.

Example:

In this example, the transition LUN is "fc_sanboot_lun" and the SCSI devices for this LUN are "/dev/sda", "/dev/sdb" and "/dev/sdc":

```
[root@IBMX3550-229-5 ~]# sanlun lun show
```

controller(7mode)/ vserver(Cmode)	lun-pathname	device filename	host adapter	protocol	lun size	mode
fas3270-229-146	/vol/fc_sanboot_lun_vol/fc_sanboot_lun	/dev/sda	host0	FCP	50g	7
fas3270-229-146	/vol/fc_sanboot_lun_vol/fc_sanboot_lun	/dev/sdb	host0	FCP	50g	7
fas3270-229-146	/vol/fc_sanboot_lun_vol/fc_sanboot_lun	/dev/sdc	host1	FCP	50g	7

3. From the OS Device ID column of the Inventory Assessment Workbook's "SAN Host LUNs" tab, identify the DMMP device name for the transition LUN.

Alternatively, you can use the "multipath -ll" command to identify the DMMP device name for the LUN on the host. The DMMP device name is either a unique device handle ID (WWID) or an alias.

Example:

In this example, the DMMP device name for SCSI device "/dev/sda" is "360a9800037534562572b453855496b47":

```
[root@IBMX3550-229-5 ~]# multipath -ll
360a9800037534562572b453855496b47 dm-0 NETAPP,LUN
size=50G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|+- policy='round-robin 0' prio=50 status=active
|  '- 1:0:0:0 sdc 8:32 active ready running
+- policy='round-robin 0' prio=10 status=enabled
|  '- 0:0:1:0 sdb 8:16 active ready running
|  '- 0:0:0:0 sda 8:0 active ready running
```

4. Verify that the DMMP device in "/dev/mapper" directory by using "ls /dev/mapper/ <DMMP device name>" command. If you cannot locate the device in the directory then DMMP device might be using an alias or user-friendly name.

Example:

In this example, "ls /dev/mapper/360a9800037534562572b453855496b47" and DMMP devices are not using either the "user_friendly_name" or an "alias":

```
[root@IBMX3550-229-5 ~]# ls /dev/mapper/360a9800037534562572b453855496b47
/dev/mapper/360a9800037534562572b453855496b47
```

5. From the "SAN Host LVMS" tab in the Inventory Assessment Workbook, determine if the DMMP device is part of an LVM.

Alternatively, you can determine whether the DMMP device is part of an LVM by using the "blkid"

command. Look at the TYPE value of the DMMP device. If the TYPE is "LVM2_member", the DMMP is part of an LVM.

Example:

In this example, DMMP device "360a9800037534562572b453855496b47" is partitioned into p1 and p2. The partition p2 is part of the LVM and p1 has the "ext4" file system:

```
[root@IBMX3550-229-5 ~]# blkid
/dev/mapper/vg_ibmx35502295-lv_root: UUID="13b11b3c-b3bf-46a1-8943-02c029e7a599" TYPE="ext4"
/dev/mapper/360a9800037534562572b453855496b47p1: UUID="8e1279b9-3beb-4c01-9ea7-c31bee157fb8" TYPE="ext4"
/dev/mapper/360a9800037534562572b453855496b47p2: UUID="ZJCwqg-nYep-otof-abUj-jqGt-59pF-k7NpJd" TYPE="LVM2_member"
/dev/mapper/vg_ibmx35502295-lv_swap: UUID="b0b288a2-2f03-4431-8e60-d21c7d8a04b2" TYPE="swap"
```

6. See the "SAN Host LVMs" tab in the Inventory Assessment Workbook to locate the LVM details, including LV names, VG name and PV devices.

Alternatively, on the host, you can use the "lvmdisplay -m" command to identify the LV information for the DMMP device.

Example:

In this example, DMMP device "360a9800037534562572b453855496b47" has 2 LVs: "/dev/vg_ibmx35502295/lv_root" and "/dev/vg_ibmx35502295/lv_swap"

```

[root@IBMx3550-229-5 ~]# lvsdisplay -m
--- Logical volume ---
LV Path                /dev/vg_ibmx35502295/lv_root
LV Name                lv_root
VG Name                vg_ibmx35502295
LV UUID                LuPcLv-5zf3-4V6S-PAMk-ukeL-AdrO-RUTgcj
LV Write Access        read/write
LV Creation host, time IBMx3550-229-5.gdl.englab.netapp.com, 2014-09-03 03:14:41 -0400
LV Status              available
# open                 1
LV Size                44.51 GiB
Current LE             11394
Segments              1
Allocation             inherit
Read ahead sectors     auto
- currently set to    256
Block device          253:3

--- Segments ---
Logical extent 0 to 11393:
  Type                linear
  Physical volume      /dev/mapper/360a9800037534562572b453855496b47p2
  Physical extents     0 to 11393

--- Logical volume ---
LV Path                /dev/vg_ibmx35502295/lv_swap
LV Name                lv_swap
VG Name                vg_ibmx35502295
LV UUID                IVIP7d-OKyf-h6RI-ZuFb-T3PB-l4TX-v5xBFu
LV Write Access        read/write
LV Creation host, time IBMx3550-229-5.gdl.englab.netapp.com, 2014-09-03 03:14:43 -0400
LV Status              available
# open                 1
LV Size                5.00 GiB
Current LE             1280
Segments              1
Allocation             inherit
Read ahead sectors     auto
- currently set to    256
Block device          253:4

--- Segments ---
Logical extent 0 to 1279:
  Type                linear
  Physical volume      /dev/mapper/360a9800037534562572b453855496b47p2
  Physical extents     11394 to 12673

```

7. From the “SCSI Device ID / Device Name” and “Mount” columns in the Inventory Assessment Workbook’s “SAN Host Filesystems” tab, identify the DMMP devices or LVM names and mount directory on which the RHEL 6 OS “/boot” and “/” (root) partitions are installed.

Alternatively, on the host, you can use either the “df -h” command or “mount” command.

Example:

The “/boot” partition is created on DMMP device “/dev/mapper/ and 360a9800037534562572b453855496b47p1” is mounted on directory “/boot”. The “p1” designation represents the partition on the DMMP device.

The “/” (root) partition is located on Logical Volume /dev/vg_ibmx35502295/lv_root”.

8. From the “Filesystem” column of the Inventory Assessment Workbook’s “SAN Host Filesystems” tab, identify the file system configured on DMMP device.

Alternatively, you can identify the file system using “blkid” command. Check the TYPE value against that of the DMMP and LV devices.

9. Check the “/etc/fstab” entries to obtain mount points details of “/” and “/boot” partitions.

Example:

In this example, LVMs “/dev/vg_ibmx35502295/lv_root” and “/dev/vg_ibmx35502295/lv_swap” are mounted on “/” and “swap” respectively:

```
[root@IBMX3550-229-5 ~]# cat /etc/fstab
#
# /etc/fstab
# Created by anaconda on Wed Sep  3 03:18:35 2014
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
/dev/mapper/vg_ibmx35502295-lv_root / ext4 defaults 1 1
UUID=8e1279b9-3beb-4c01-9ea7-c31bee157fb8 /boot ext4 defaults 1 2
/dev/mapper/vg_ibmx35502295-lv_swap swap swap defaults 0 0
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
```

10. If “/” and “swap” partitions are mounted using LV names, no pretransition changes are required in the “/etc/fstab” file.
11. If the “/boot” partition is installed on the DMMP device, check how it is referenced to mount in the “/etc/fstab” at boot time.

Example:

The “/boot” partition is reference via its file system UUID number: 8e1279b9-3beb-4c01-9ea7-c31bee157fb8, and this UUID number will be used at boot time to mount “/boot” partition.

12. If the “/boot” partition is mounted using the file system UUID obtained from the “blkid” command output, then no pretransition changes are required.

1.16.2 Before entering cutover phase

1. For FC configurations, ensure that fabric connectivity and zoning to clustered Data ONTAP controllers is established. For iSCSI configurations, ensure that the iSCSI sessions to clustered Data ONTAP controllers are discovered and logged in.
2. Shutdown the host.

1.16.3 Posttransition

After completing the transition to clustered Data ONTAP, complete the following steps:

1. Power on the host.
2. Enter HBA BIOS settings mode.
3. Choose “Rescan” to discover the clustered Data ONTAP SANBoot lun.

4. Clear the 7-Mode Boot LUN ID.
5. Add the clustered Data ONTAP Boot LUN ID in HBA BIOS.
6. Exit HBA BIOS and reboot the host.
7. Wait for the OS boot to complete.
8. Execute the “sanlun lun show” command to verify the LUN:

```
[root@IBMX3550-229-5 ~]# sanlun lun show
controller(7mode)/
vserver(Cmode)      lun-pathname      device
                    filename      host
                    adapter      protocol      lun
                    size      mode
-----
vs_7MIT             /vol/fc_sanboot_lun_vol_cdot/fc_sanboot_lun /dev/sdb      host0      FCP      50g      C
vs_7MIT             /vol/fc_sanboot_lun_vol_cdot/fc_sanboot_lun /dev/sdg      host1      FCP      50g      C
vs_7MIT             /vol/fc_sanboot_lun_vol_cdot/fc_sanboot_lun /dev/sdh      host1      FCP      50g      C
vs_7MIT             /vol/fc_sanboot_lun_vol_cdot/fc_sanboot_lun /dev/sdf      host1      FCP      50g      C
vs_7MIT             /vol/fc_sanboot_lun_vol_cdot/fc_sanboot_lun /dev/sdd      host0      FCP      50g      C
vs_7MIT             /vol/fc_sanboot_lun_vol_cdot/fc_sanboot_lun /dev/sdc      host0      FCP      50g      C
vs_7MIT             /vol/fc_sanboot_lun_vol_cdot/fc_sanboot_lun /dev/sde      host1      FCP      50g      C
vs_7MIT             /vol/fc_sanboot_lun_vol_cdot/fc_sanboot_lun /dev/sda      host0      FCP      50g      C
```

9. Execute “multipath -ll” to verify the DMMP device:

```
[root@IBMX3550-229-5 ~]# multipath -ll
3600a098051764937303f4479515a7454 dm-0 NETAPP,LUN C-Mode
size=50G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|  |- 0:0:2:0 sdc 8:32 active ready running
|  |- 1:0:4:0 sdh 8:112 active ready running
|  |- 1:0:3:0 sdg 8:96 active ready running
|  `-- 0:0:1:0 sdb 8:16 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
|  |- 0:0:0:0 sda 8:0 active ready running
|  |- 1:0:0:0 sde 8:64 active ready running
|  |- 0:0:5:0 sdd 8:48 active ready running
|  `-- 1:0:2:0 sdf 8:80 active ready running
```

1.17 Protocol: iSCSI

1.17.1 Pretransition

1. On the 7-Mode controller console, identify the 7-Mode LUN name on which the RHEL 6 OS is installed by using the “LUN Name” column of the Inventory Assessment Workbook’s “LUN” tab.

Alternatively, on the host, use the 7-Mode “lun show” command to list the LUNs on the controller.

2. From the “SCSI Device ID” column of the Inventory Assessment Workbook’s “SAN Host LUNs” tab, identify the SCSI device names for the LUN being transitioned.

Alternatively, you can use the “sanlun lun show” command on the host to identify the SCSI device name for a LUN.

Example:

In this example, the transition LUN is “lun_os” and the SCSI devices for this LUN are “/dev/sdc” and “/dev/sdb”:

```
[root@IBMX3550M3-229-167 ~]# sanlun lun show
controller(7mode)/
vserver(Cmode)      lun-pathname      device
                    filename      host
                    adapter      protocol      lun
                    size      mode
-----
fas3270-229-146     /vol/lun_os_vol/lun_os /dev/sdc      host6      iSCSI     50g      7
fas3270-229-146     /vol/lun_os_vol/lun_os /dev/sdb      host5      iSCSI     50g      7
```

3. From the “OS Device ID” column of the Inventory Assessment Workbook’s “SAN Host LUNs” tab, identify the DMMP device name for the transition LUN.

Alternatively, you can use the “multipath -ll” command to identify the DMMP device name for the LUN on the host.

Example:

In this example, the DMMP device name for SCSI device “/dev/sdb” is “360a9800037534562572b453855496b45”

```
360a9800037534562572b453855496b45 dm-2 NETAPP,LUN
size=50G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=2 status=active
   |-- 5:0:0:0 sdb 8:16 active ready running
   `-- 6:0:0:0 sdc 8:32 active ready running
```

4. Verify that the DMMP device exists in the “/dev/mapper” directory by running the “ls /dev/mapper/ <DMMP device name>” command. If the DMMP device is not displayed, the device might be using an alias or a user-friendly name.

Example:

In this example, “ls /dev/mapper/360a9800037534562572b453855496b45”

```
[root@IBMX3550M3-229-167 ~]# ls /dev/mapper/360a9800037534562572b453855496b45
/dev/mapper/360a9800037534562572b453855496b45
```

5. From the “SAN Host LVMS” tab in the Inventory Assessment Workbook, determine if DMMP device is part of an LVM.

Alternatively, you can determine whether the DMMP device is part of an LVM by using the “blkid” command. Look at the TYPE value of the DMMP device. If the TYPE is “LVM2_member”, the DMMP is part of an LVM.

Example:

In this example, DMMP device “360a9800037534562572b453855496b45” is part of an LVM:

```
[root@IBMX3550M3-229-167 ~]# blkid
/dev/mapper/vg_ibmx3550m3229-LogVol00: UUID="7a1e569d-5176-48a9-b638-569b027973d9" TYPE="ext4"
/dev/sda1: UUID="2f81d26f-78c3-4b53-9f58-aa31c139ffb4" TYPE="ext4"
/dev/mapper/3600605b00425863016fe22f30862c5b1p1: UUID="2f81d26f-78c3-4b53-9f58-aa31c139ffb4" TYPE="ext4"
/dev/mapper/360a9800037534562572b453855496b45p1: UUID="p4aPI9-Mant-6rjF-JXyi-ireJ-EVTZ-zvpwP" TYPE="LVM2_member"
/dev/mapper/vg_ibmx3550m3229-LogVol01: UUID="9e423dcb-f2f5-4d23-a857-93a075e5e04c" TYPE="swap"
```

6. See the “SAN Host LVMS” tab in the Inventory Assessment Workbook to get the LVM the following details: LV names, VG name and PV devices.

Alternatively, you can use the “lvdisplay -m” command on the host to determine the LV information for the DMMP device.

Example:

In this example, DMMP device “360a9800037534562572b453855496b45” has two LVs: “/dev/vg_ibmx3550m3229/LogVol01” and “/dev/vg_ibmx3550m3229/LogVol00”

```
[root@IBMX3550M3-229-167 ~]# lvdisplay -m
--- Logical volume ---
LV Path                /dev/vg_ibmx3550m3229/LogVol101
LV Name                 LogVol101
VG Name                 vg_ibmx3550m3229
LV UUID                 kpVOK2-1IuG-SBtd-id0A-WqyB-9zXD-sVxWza
LV Write Access         read/write
LV Creation host, time  IBMx3550M3-229-167.gdl.englab.netapp.com, 2014-08-12 14:27:11 +0530
LV Status                available
# open                  1
LV Size                  4.00 GiB
Current LE               32
Segments                 1
Allocation               inherit
Read ahead sectors       auto
- currently set to       256
Block device             253:4

--- Segments ---
Logical extent 0 to 31:
  Type                    linear
  Physical volume          /dev/mapper/360a9800037534562572b453855496b45p1
  Physical extents         0 to 31

--- Logical volume ---
LV Path                /dev/vg_ibmx3550m3229/LogVol100
LV Name                 LogVol100
VG Name                 vg_ibmx3550m3229
LV UUID                 MiUmWt-HXKC-ntOQ-sA8W-A6Q2-KDWh-kXRE17
LV Write Access         read/write
LV Creation host, time  IBMx3550M3-229-167.gdl.englab.netapp.com, 2014-08-12 14:27:11 +0530
LV Status                available
# open                  1
LV Size                  19.62 GiB
Current LE               157
Segments                 1
Allocation               inherit
Read ahead sectors       auto
- currently set to       256
Block device             253:5

--- Segments ---
Logical extent 0 to 156:
  Type                    linear
  Physical volume          /dev/mapper/360a9800037534562572b453855496b45p1
  Physical extents         32 to 188
```

7. From the “SCSI Device ID / Device Name” and “Mount” columns in the Inventory Assessment Workbook’s “SAN Host Filesystems” tab, identify the DMMP devices or LVM names and mount directory on which the RHEL 6 OS “/boot” and “/” (root) partitions are installed.

Alternatively, on the host, you can use either the “df -h” command or “mount” command.

Example:

The “/boot” partition is created on DMMP device “/dev/mapper/ and 360a9800037534562572b453855496b45p1” is mounted on directory “/boot”. The “p1” designation represents the partition on the DMMP device.

The “/” (root) partition is installed on Logical Volume /dev/vg_ibmx3550m3229/LogVol100”.

8. From the “Filesystem” column of the Inventory Assessment Workbook’s “SAN Host Filesystems” tab, identify the file system configured on DMMP.
Alternatively, you can identify the file system using “blkid” command. Check the TYPE value against that of the DMMP and LV devices.
9. Check the “/etc/fstab” entries to obtain mount point details of “/” and “/boot” partitions.

Example:

In this example, LVMs “/dev/vg_ibmx3550m3229/LogVol01” and “/dev/vg_ibmx3550m3229/LogVol00” are mounted on “swap” and “/” respectively:

```
[root@IBMX3550M3-229-167 ~]# cat /etc/fstab
#
# /etc/fstab
# Created by anaconda on Tue Aug 12 14:27:43 2014
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
/dev/mapper/vg_ibmx3550m3229-LogVol00 /                ext4    defaults,_rnetdev 1 1
UUID=2f81d26f-78c3-4b53-9f58-aa31c139ffb4 /boot          ext4    defaults          1 2
/dev/mapper/vg_ibmx3550m3229-LogVol01 swap           swap    defaults,_netdev 0 0
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
```

If the “/” and “swap” partitions are mounted using LV names, no pretransition changes are required in the “/etc/fstab” file.

10. If the “/boot” partition is installed on DMMP device, check how it is referenced to mount in “/etc/fstab” at boot time.

Example:

The “/boot” partition is referenced by the file system UUID number: 2f81d26f-4b53-9f58-aa31c139ffb4. This UUID number is used at boot time to mount the “/boot” partition.

11. 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.
12. For iSCSI SAN booted hosts, you can obtain iSCSI sessions for 7-Mode controllers from the Inventory Assessment Workbook. See the “Target IPs (iSCSI)” column in the “SAN Host HBAs” tab. You can also see the “iSCSI SAN Interfaces” tab to identify the iSCSI IQN number and IPs that are configured on 7-Mode controllers.

Alternatively, execute “iscsiadm -m sessions” command on the host.

Example:

In this example, iSCSI sessions 1 and 2 are from 7-Mode controllers and 3, 4, 5 and 6 are from clustered Data ONTAP controllers:

```
root@IBMX3550M3-229-167 ~]# iscsiadm -m session
cp: [1] 10.226.229.241:3260,1000 iqn.1992-08.com.netapp:sn.1574168453
cp: [2] 10.226.228.155:3260,1001 iqn.1992-08.com.netapp:sn.1574168453
cp: [3] 10.226.228.99:3260,1081 iqn.1992-08.com.netapp:sn.ac1f53a61d3c11e4b90700a0983afb96:vs.4
cp: [4] 10.226.228.98:3260,1080 iqn.1992-08.com.netapp:sn.ac1f53a61d3c11e4b90700a0983afb96:vs.4
cp: [5] 10.226.228.97:3260,1079 iqn.1992-08.com.netapp:sn.ac1f53a61d3c11e4b90700a0983afb96:vs.4
cp: [6] 10.226.228.96:3260,1078 iqn.1992-08.com.netapp:sn.ac1f53a61d3c11e4b90700a0983afb96:vs.4
```

13. For iSCSI SAN booted host “/boot/grub/grub.conf” file should be edited to create a new kernel command line entry that includes the clustered Data ONTAP controller’s IQN number and iSCSI session information.

See the “GRUB Configuration” and “GRUB Device Map” tabs in the Inventory Assessment Workbook to locate grub information.

Example:

This example illustrates the “/boot/grub/grub.conf” file before editing. The kernel command line has

the 7-Mode controller's IQN number and iSCSI session information.

```
# grub.conf generated by anaconda
#
# Note that you do not have to rerun grub after making changes to this file
# NOTICE: You have a /boot partition. This means that
#         all kernel and initrd paths are relative to /boot/, eg.
#         root (hd0,0)
#         kernel /vmlinuz-version ro root=/dev/mapper/vg_ibmx3550m3229-LogVol100
#         initrd /initrd-[generic-]version.img
#boot=/dev/sda
default=0
timeout=5
splashimage=(hd0,0)/grub/splash.xpm.gz
hiddenmenu
title Red Hat Enterprise Linux (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:@10.226.229.241::3260::iqn.1992-08.com.netapp:sn.1574168453 LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx3550m3229/LogVol01 rd_LVM_LV=vg_ibmx3550m3229/LogVol00 rd_NO_MD netroot=iscsi:@10.226.228.155::3260::iqn.1992-08.com.netapp:sn.1574168453 iscsi_initiator=iqn.1994-05.com.redhat:229.167 crashkernel=auto SYSFONT=latarcyrheb-sun16 KEYBOARDTYPE=pc KEYTABLE=us ip=eth0:dhcp rd_NO_DM rhgb quiet
    initrd /initramfs-2.6.32-431.el6.x86_64.img
```

Example 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 (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.ac1f53a61d3c11e4b90700a0983afb96:vs.4 LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx3550m3229/LogVol01 rd_LVM_LV=vg_ibmx3550m3229/LogVol00 rd_NO_MD netroot=iscsi:@10.226.228.98::3260::iqn.1992-08.com.netapp:sn.ac1f53a61d3c11e4b90700a0983afb96:vs.4 netroot=iscsi:@10.226.228.97::3260::iqn.1992-08.com.netapp:sn.ac1f53a61d3c11e4b90700a0983afb96:vs.4 netroot=iscsi:@10.226.228.96::3260::iqn.1992-08.com.netapp:sn.ac1f53a61d3c11e4b90700a0983afb96:vs.4 iscsi_initiator=iqn.1994-05.com.redhat:229.167 crashkernel=auto SYSFONT=latarcyrheb-sun16 KEYBOARDTYPE=pc KEYTABLE=us ip=eth0:dhcp rd_NO_DM rhgb quiet
    initrd /initramfs-2.6.32-431.el6.x86_64.img
```

14. Take the backup of the initrd image:

```
[root@IBMX3550M3-229-167 ~]# cp /boot/initramfs-2.6.32-431.el6.x86_64.img /boot/initramfs-2.6.32-431.el6.x86_64.img_backup
```

15. Update the old 7-Mode kernel line in “/boot/grub/grub.conf” with the backup initrd image name:

```

title Red Hat Enterprise Linux (2.6.32-431.el6.x86_64) 7Mode
    root (hd0,0)
    kernel /vmlinuz-2.6.32-431.el6.x86_64 ro root=/dev/mapper/vg_ibmx3550m32
29-LogVol100 ifname=eth0:5c:f3:fc:ba:46:d8 rd_NO_LUKS netroot=iscsi:@10.226.229.2
41::3260::iqn.1992-08.com.netapp:sn.1574168453 LANG=en_US.UTF-8 rd_LVM_LV=v
vg_ibmx3550m3229/LogVol101 rd_LVM_LV=v
vg_ibmx3550m3229/LogVol100 rd_NO_MD netroot=iscsi:@
10.226.228.155::3260::iqn.1992-08.com.netapp:sn.1574168453 iscsi_initiator=iqn.1
994-05.com.redhat:229.167 crashkernel=auto SYSFONT=latacyrheb-sun16 KEYBOARDTY
PE=pc KEYTABLE=us ip=eth0:dhcp rd_NO_DM rhgb quiet
    initrd /initramfs-2.6.32-431.el6.x86_64.img_backup

```

16. Make sure the clustered Data ONTAP kernel line in the “/boot/grub/grub.conf” file is appended with “rdloaddriver=scsi_dh_alua”:

```

title Red Hat Enterprise Linux (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_ibmx3550m32
29-LogVol100 ifname=eth0:5c:f3:fc:ba:46:d8 rd_NO_LUKS netroot=iscsi:@10.226.228.9
9::3260::iqn.1992-08.com.netapp:sn.ac1f53a61d3c11e4b90700a0983afb96:vs.4 LANG=en
_US.UTF-8 rd_LVM_LV=v
vg_ibmx3550m3229/LogVol101 rd_LVM_LV=v
vg_ibmx3550m3229/LogVol100 rd_NO_MD netroot=iscsi:@10.226.228.98::3260::iqn.1992-08.com.netapp:sn.ac1f53a
61d3c11e4b90700a0983afb96:vs.4 netroot=iscsi:@10.226.228.97::3260::iqn.1992-08.c
om.netapp:sn.ac1f53a61d3c11e4b90700a0983afb96:vs.4 netroot=iscsi:@10.226.228.96:
:3260::iqn.1992-08.com.netapp:sn.ac1f53a61d3c11e4b90700a0983afb96:vs.4 iscsi_ini
tiator=iqn.1994-05.com.redhat:229.167 crashkernel=auto SYSFONT=latacyrheb-sun16
    KEYBOARDTYPE=pc KEYTABLE=us ip=eth0:dhcp rd_NO_DM rhgb quiet rdloaddriver=scsi
_dh_alua
    initrd /initramfs-2.6.32-431.el6.x86_64.img

```

17. If “/etc/multipath.conf” or “/boot/grub/grub.conf” files are updated then update the kernel initial RAM disk (initramfs). The initramfs must be recreated so that the new clustered Data ONTAP IQN number and iSCSI sessions are referenced and so the host establishes an iSCSI connection with clustered Data ONTAP controllers at boot time.
18. Back up the existing initramfs.

Example:

```

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

```

19. Recreate the initrd image by using the “dracut -force --add multipath --verbose” command.

1.17.2 Before entering cutover phase

1. For FC configurations, ensure that fabric connectivity and zoning to clustered Data ONTAP controllers is established. For iSCSI configurations, ensure that the iSCSI sessions to clustered Data ONTAP controllers are discovered and logged in.
2. Shut down the host.

1.17.3 Posttransition

1. After completing the transition to clustered Data ONTAP, power on the host.
2. Execute “sanlun lun show” command to verify the LUN:


```
[root@IBMx3550M3-229-167 ~]# sanlun lun show
controller(7mode)/
vserver(Cmode)      lun-pathname      device
                    filename      host
                    adapter      protocol      lun
                    size      mode
-----
vs_7MTT             /vol/cdot_lun_os_vol/lun_os /dev/sde      host8      iSCSI      50g      C
vs_7MTT             /vol/cdot_lun_os_vol/lun_os /dev/sdd      host7      iSCSI      50g      C
vs_7MTT             /vol/cdot_lun_os_vol/lun_os /dev/sdc      host6      iSCSI      50g      C
vs_7MTT             /vol/cdot_lun_os_vol/lun_os /dev/sdb      host5      iSCSI      50g      C
```

3. Execute “multipath -ll” to verify DMMP device.

Windows host transition

1.18 Pretransition

1. Install the 7MTT Tool. See the documentation for more information.
2. Identify the LUN serial numbers, LUN ID's, and the corresponding Windows physical disk numbers of the LUNs being transitioned. On systems running Data **ONTAPDSM** multipathing solution, this can be achieved in two ways:
 - By using the Data ONTAPDSM Management Extension Snap-In (accessible through Server Manager or Computer Management under Storage). See the following screenshot:



- Alternatively, this can be achieved using the get-sandisk Windows PowerShell cmdlet that is included as part of Data ONTAP DSM. See the following screenshot:

```

Administrator: Windows PowerShell
PS C:\Users\administrator.WSHE> Get-sandisk

DiskId DriveLetters SerialNumber Size LbPolicy PathCount LUN LunType Storage LunPath Mode Reserved
-----
Disk0 Y: 2F1JUS8VK4BR 1027 M DLQD 4 100 windows_2008 wshe3170c-09 /vol/vol125/test1_7mtt SingleImage No
Disk1 Y: 2F1JUS8VK4BT 1027 M DLQD 4 101 windows_2008 wshe3170c-09 /vol/vol126/test2_7mtt SingleImage No
Disk2 T2-KS7DYQDEVP 30 G DLQD 4 51 windows_2008 vs_smb_vm /vol/vol1_smb_vm_2k12_w2k12boot/smb_test C-Mode Yes
Disk3 T2-KS7DYQDEVR 30 G DLQD 4 52 windows_2008 vs_smb_vm /vol/vol1_smb_vm_intervolume/csv_intervol C-Mode Yes
Disk4 D2-K0SDUT2IH 10 G DLQD 8 0 windows_2008 vs_wshex3650_88_89_fc /vol/vs_wshex3650_88_fc_vol_1/vs_wshex3650_88_fc_lun C-Mode No
Disk5 I: D2-MAJQYODEF 10 G DLQD 8 1 windows_2008 vs_wshex3650_88_89_fc /vol/vs_wshex3650_88_fc_vol_2/vs_wshex3650_88_fc_lun C-Mode No
Disk6 D2-KS7DYQDEH 10 G DLQD 8 2 windows_2008 vs_wshex3650_88_89_fc /vol/vs_wshex3650_88_fc_vol_3/vs_wshex3650_88_fc_lun C-Mode No
Disk7 N: D2-K9+DYODEJ 10 G DLQD 8 3 windows_2008 vs_wshex3650_88_89_fc /vol/vs_wshex3650_88_fc_vol_4/vs_wshex3650_88_fc_lun C-Mode No
Disk8 M: D2-K0SDUT2II 10 G DLQD 8 4 windows_2008 vs_wshex3650_88_89_fc /vol/vs_wshex3650_89_fc_vol_1/vs_wshex3650_89_fc_lun C-Mode Yes
Disk9 C: D2-K0SDUT2H6 120 G DLQD 8 0 windows_2008 boot_vserver /vol/boot_vol/w2k12_88 C-Mode No
Disk10 D2-K0SDUT2HY 2055 M DLQD 8 25 windows_2008 boot_vserver /vol/boot_vol/w2k12_quo C-Mode Yes
Disk11 T2-K0SDUT3AE 80 G DLQD 4 50 hyper_v boot_vserver /vol/boot_vol/w2k12_vm_vhd_clus C-Mode Yes
Disk12 F: D2-K0SDUT2II 10 G DLQD 8 0 windows_2008 vs_wshex3650_88_89_fc /vol/vs_wshex3650_88_fc_vol_1/vs_wshex3650_88_fc_lun C-Mode Yes
Disk13 D2-MAJQYODEH 10 G DLQD 8 1 windows_2008 vs_wshex3650_88_89_fc /vol/vs_wshex3650_88_fc_vol_2/vs_wshex3650_88_fc_lun C-Mode No
Disk14 D2-K9+DYODEK 10 G DLQD 8 2 windows_2008 vs_wshex3650_88_89_fc /vol/vs_wshex3650_88_fc_vol_4/vs_wshex3650_88_fc_lun C-Mode No
Disk15 D2-KS7DYODEJ 10 G DLQD 8 3 windows_2008 vs_wshex3650_88_89_fc /vol/vs_wshex3650_88_fc_vol_3/vs_wshex3650_88_fc_lun C-Mode No
Disk16 L: D2-MAJQYODEI 10 G DLQD 8 4 windows_2008 vs_wshex3650_88_89_fc /vol/vs_wshex3650_89_fc_vol_2/vs_wshex3650_89_fc_lun C-Mode Yes
Disk17 K: D2-K0SDUT2IW 11 G DLQD 8 0 windows_2008 vs_wshex3650_88_89_iscsi /vol/vs_wshex3650_88_iscsi_vol_1/vs_wshex3650_88_iscsi_lun C-Mode Yes
Disk18 J: D2-MAJQYODEW 12 G DLQD 8 1 windows_2008 vs_wshex3650_88_89_iscsi /vol/vs_wshex3650_88_iscsi_vol_2/vs_wshex3650_88_iscsi_lun C-Mode No
Disk19 D2-K9+DYODE/ 10 G DLQD 8 2 windows_2008 vs_wshex3650_88_89_iscsi /vol/vs_wshex3650_88_iscsi_vol_4/vs_wshex3650_88_iscsi_lun C-Mode No
Disk20 D2-MAJQYODEX 10 G DLQD 8 4 windows_2008 vs_wshex3650_88_89_iscsi /vol/vs_wshex3650_89_iscsi_vol_2/vs_wshex3650_89_iscsi_lun C-Mode No
Disk21 T2-MAJQYODE1 3074 M DLQD 4 250 windows_2008 vs_wshex3650_88_89_iscsi /vol/vol10/iscsi_test C-Mode No

PS C:\Users\administrator.WSHE>

```

See the *Data ONTAP DSM Installation and Administration Guide* for more details on Data ONTAPDSM Management Extension Snap-In or Windows PowerShell cmdlets.

On systems running **MSDSM** as the multipathing solution, use the **Inventory Collect Tool** to gather the LUN serial numbers, LUN IDs and the corresponding Windows physical disk numbers of the LUNs being transitioned. Inventory Collect Tool (ICT) is a standalone 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. The tool generates an Inventory Assessment Workbook and an Inventory Report XML file that contain configuration details of the storage and host systems. The LUN ID, LUN serial number and the corresponding serial number is captured under the SAN Host LUNs tab. See the Inventory Collect Tool Host and Storage Information Collection Guide for more details on ICT.

1. If the LUNs being transitioned are FC/FCoE LUNs, create or modify fabric zoning to ensure the LUNs are visible to the host after the transition completes. See the *SAN Administration Guide* for additional details on zoning.
2. If the LUNs being transitioned are iSCSI LUNs, create iSCSI sessions that connect to the clustered Data ONTAP controller so that the LUNs are visible after transition. See the *Microsoft iSCSI Software Initiator Users Guide* for details on creating iSCSI sessions.


1.19 Transition

1. Run the 7MTT Tool to transfer your volumes from Data ONTAP operating in 7-Mode to clustered Data ONTAP. Follow the instructions for any warnings displayed in the pre-checks.
2. **Before starting the cutover phase**, do the following:
 - On systems running ONTAPDSM as the multipathing solution, ensure the version of ONTAP DSM installed on the server is supported for the version of Data ONTAP running


on the target clustered Data ONTAP controller.

- On systems running MSDSM as the multipathing solution, ensure the version of Windows Host Utilities installed on the server is supported for the version of Data ONTAP running on the target clustered Data ONTAP controller.

3. Take offline the disks which must be transitioned from the Disk Manager. See the following screenshot:



4. If the host is booted from SAN and if the boot lun is being transitioned, shut down the boot host. Steps 5 and 6 can be skipped in this case.
5. If the host is clustered, offline the clustered disks, including the quorum disk, using Failover Cluster Manager.
6. If the host you are transitioning is running Windows Server 2003 and you need to migrate the quorum device, you should stop the cluster services on all cluster nodes.




7. If you are transitioning LUNs on a server that has Hyper-V enabled, do one of the following:
 - If you are transitioning a Windows guest OS, see the [Windows host transition steps](#).
 - If you are transitioning a Linux guest OS, see the [Linux host transition steps](#).
8. If you are transitioning LUNs on a server that has Hyper-V enabled, and the boot device of the guest OS resides on a Data ONTAP LUN that is being transitioned, shut down the guest operating system before taking the corresponding disk offline on the parent system.

1.20 Posttransition

1. Invoke the transition CLI “transition export lunmap” command to generate mapping of the 7-Mode LUNs to the transitioned LUNs for the transition session.
2. If the Windows host is SAN-booted and the boot LUN was transitioned, power on the host. Update the FC BIOS to enable the system to boot from the LUN on the clustered Data ONTAP controller. See the HBA vendor documentation for more information.
3. On the Windows host, rescan the disks from the Disk Manager.
4. Verify if the transitioned LUNs are showing up on the Host
 - On systems running ONTAPDSM as the multipathing solution, use Data ONTAPDSM Management Extension Snap-In or the get-sandisk Windows PowerShell cmdlet to obtain the LUN serial numbers, LUN IDs and the corresponding Windows physical disk numbers of the LUNs mapped to the host. When using this data along with the LUN map output and the data collected in pretransition state, you can determine whether the LUNs have transitioned successfully. Note whether the Physical disk numbers of the transitioned LUNs have changed.
 - On systems running MSDSM as the multipathing solution, use the Inventory Collect Tool to gather the LUN serial numbers, LUN ID's and the corresponding Windows physical disk numbers of the LUNs mapped to the host. The LUN ID, LUN serial number and the corresponding serial number is captured under the SAN Host LUNs tab. Refer to the

Inventory Collect Tool Host and Storage Information Collection Guide for more details on ICT. When using this data along with the LUN map output and the data collected in pretransition state, you can determine whether the LUNs have transitioned successfully. Note whether the Physical disk numbers of the transitioned LUNs have changed.

5. Bring the disks online from the Disk Manager.
6. If Windows Failover Cluster is configured on the host, bring the disks online from the Failover Cluster Manager.
7. If the host you are transitioning is running Windows Server 2003 and you have migrated the quorum device, start the cluster services on all the cluster nodes.
8. If Hyper-V is enabled on the host and pass-through devices configured to the VMs, modify the settings from Hyper-V Manager in case the Physical disk number of the LUN corresponding to the pass-through device has changed after the transition.



9. If you are using Windows 2003, upgrade to windows 2008 prior to transitioning to the clustered Data ONTAP controller. Windows 2003 is not supported by Data ONTAP 8.3 and later. You must ensure that LUN types are set correctly to prevent misalignment issues.

1.21 Exceptions and known issues

1. For transitioning Hyper-V VMs with only VHD or VHDX file types, Storage live-migration can be used instead of 7MTT. Refer to the Microsoft documentation for more details on Hyper-V storage live migration.
2. If you used the same igroup name on both nodes of the 7-Mode controller, the transition tool might fail to resolve igroup conflict. See Bugs Online ID 804962. [Bugs Online - support.netapp.com/NOW/cgi-bin/bol/](https://support.netapp.com/NOW/cgi-bin/bol/).

Copyright information

Copyright © 1994–2015 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, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

How to send comments about documentation and receive update notification

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277