



# **VMware ESXi host remediation**

## **ONTAP 7-Mode Transition**

NetApp  
February 11, 2024

This PDF was generated from [https://docs.netapp.com/us-en/ontap-7mode-transition/san-host/concept\\_esxi\\_versions\\_and\\_features\\_supported\\_for\\_san\\_transitions\\_using\\_7mtt.html](https://docs.netapp.com/us-en/ontap-7mode-transition/san-host/concept_esxi_versions_and_features_supported_for_san_transitions_using_7mtt.html) on February 11, 2024. Always check docs.netapp.com for the latest.

# Table of Contents

- VMware ESXi host remediation . . . . . 1
  - ESXi versions and features supported for SAN transitions using 7MTT . . . . . 1
  - Preparing for transition of ESXi hosts . . . . . 1
  - Testing transitioned LUNs and ESXi host applications before the cutover phase of copy-based transitions . 6
  - Downtime at apply configuration (precutover) phase of ESXi host transition . . . . . 8
  - Post-transition remediation requirements for ESXi hosts. . . . . 8
  - Enabling CAW on a datastore using the ESXi CLI. . . . . 17

# VMware ESXi host remediation

If you are using the 7-Mode Transition Tool (7MTT) to move from Data ONTAP operating in 7-Mode to clustered Data ONTAP in a SAN environment, you must perform a series of steps on your VMware ESXi hosts before the Data ONTAP transition. You must power down your hosts before cutover and you must perform another series of steps after transition before you can begin servicing data.

## Related information

[Testing transitioned LUNs and ESXi host applications before the cutover phase of copy-based transitions](#)

[Post-transition remediation requirements for ESXi hosts](#)

## ESXi versions and features supported for SAN transitions using 7MTT

Only certain versions and features of ESXi are supported for SAN transitions using the 7-Mode Transition Tool (7MTT).

The following versions and features are supported as listed in the [NetApp Interoperability Matrix Tool](#)

- ESXi 5.0, 5.1, 5.5, and later

You must upgrade hosts running ESX/ESXi 4.x or earlier to ESX/ESXi 5.0 or later to transition them.

- VMFS3 and VMFS5 datastores
- SAN boot configurations
- RDM (RAW device map) devices
- All Guest OS supported in the Interoperability Matrix
- All SAN protocols (FC/FCoE/iSCSI)

## Preparing for transition of ESXi hosts

You must complete several prerequisite tasks before you use the 7-Mode Transition Tool (7MTT) to transition your ESXi hosts from Data ONTAP operating in 7-Mode to ONTAP.

### Steps

1. Configure clustered Data ONTAP as described in the [7-Mode Transition Tool Copy-Based Transition Guide](#) or the [7-Mode Transition Tool Copy-Free Transition Guide](#) based on the type of transition that you are performing.
2. Gather the following information for the ESXi hosts that you are transitioning:
  - IP address
  - Host name
  - Authentication details
3. Complete the zoning between FC or FCoE hosts and new clustered Data ONTAP nodes.

You can use the Collect and Assess feature to generate the zoning plan.

4. Use the [NetApp Interoperability Matrix Tool](#) to verify that the following are supported for transition to clustered Data ONTAP:
  - Your version of Data ONTAP operating in 7-Mode

In some cases, you might have to upgrade your version of Data ONTAP operating in 7-Mode to a 7MTT SAN compatible version. For example, Data ONTAP 7.3.7 operating in 7-Mode is not compatible for transitions using the 7MTT. If you are running this version, you must upgrade it before initiating transition.

- Your ESXi host configuration
- Your HBA driver and firmware

For iSCSI, only software initiators are supported. For FC and FCoE, only QLogic and Emulex initiators are supported. If your ESXi FC or FCoE initiator is not supported, you must upgrade to a version that is supported by clustered Data ONTAP as described in the Interoperability Matrix.

5. If configured, disable VMware high availability (HA) and Distributed Resource Scheduler (DRS).

VMware HA and DRS are not supported during transition.

## Related information

[Retaining resource pools when disabling VMware DRS clusters in the vSphere Web Client](#)

[Disabling VMware High Availability \(HA\)](#)

## What the Inventory Collect Tool is

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

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

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

## Preparing Linux guest operating systems for transition

If have 7-Mode LUNs mapped as physical compatible RDM (PTRDM) to Linux virtual machines (VMs) for the boot device, there are steps you must perform to prepare your Linux VMs for transition.

- For copy-based transitions, perform these steps before initiating the Storage Cutover operation in the 7-Mode Transition Tool (7MTT).
- For copy-free transitions, perform these steps before initiating the Export & Halt 7-Mode Systems operation in the 7MTT.

## Steps

1. Obtain the SCSI device serial numbers:

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

In the following example, 360a9800032466879362b45777447462d-part2 and 360a9800032466879362b45777447462d-part1 are SCSI device numbers:

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

2. Determine the mapping between the SCSI device serial numbers and SCSI devices/partitions:

```
# ls -l /dev/disk/by-id
```

The following example shows how the relationship mapping is displayed. The SCSI devices/partitions are displayed following the SCSI device/partition serial numbers. In this example, ../../sda, ../../sda1, and ../../sda2 are SCSI devices/partitions.

```
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
```

3. Determine the mapping between the SCSI device paths and the UUIDs:

```
ls -l /dev/disk/by-uuid
```

The following example shows how the relationship mapping is displayed. In this example, 33d43a8b-cfae-4ac4-9355-36b479cfa524 is the UUID for SCSI device/partition sda2, 603e01f8-7873-440a-9182-878abff17143 is the UUID for SCSI device/partition sdb, and c50b757b-0817-4c19-8291-0d14938f7f0f is the UUID for SCSI device/partition sda1.

```
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
```

4. Use the UUID to update the device reference in the grub boot menu.lst file by matching it with the SCSI

device path and SCSI 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"
```

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

The following example shows the menu.lst file after it has been updated:

```
# 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
```

6. Update the /etc/fstab file:

- a. Use the UUID you just retrieved to update the device reference in the /etc/fstab file.

The following example shows an /etc/fstab file with a SCSI serial number:

```
/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
```

- b. Replace the reference to the SCSI serial number with the UUID.

The following example shows an `/etc/fstab` file that has been updated to replace the SCSI serial number with the UUID:

```
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
```

## Preparing Windows guest operating systems for transition

If your Windows VMs use physical compatible RDM (PTRDM) devices, you must take the disks offline on the Windows VM before transition. You can use Disk Manager to take the disks offline.

- For copy-based transitions, perform these steps before initiating the Storage Cutover operation in the 7-Mode Transition Tool (7MTT).
- For copy-free transitions, perform these steps before initiating the Export & Halt 7-Mode Systems operation in the 7MTT.

## How to identify VM snapshots that must be removed before transition

Snapshots virtual machines (VMs) with Virtual RDM attached do not survive transition from Data ONTAP operating in 7-Mode to clustered Data ONTAP. These snapshots must be removed before transition. Snapshots of VMs with only VMFS vDisks and Physical RDM (PTRDM) do survive transition and do not have to be removed.

You can use the *Inventory Assessment Workbook* generated by the Inventory Collect Tool to identify all VMs with Virtual RDMs attached. Snapshots listed in the *Inventory Assessment Workbook* under the VM Snapshots column and the NPTRDM column with a value greater than 0 are VMs that have a Virtual RDM attached with VM snapshots.

### Deleting VM snapshots copies using vSphere Client

If you are not familiar with the ESXi CLI or if it is more convenient for your environment, you can delete Virtual Machine (VM) snapshots using vSphere Client.

- For copy-based transitions, perform these steps before initiating the Storage Cutover operation in the 7-Mode Transition Tool (7MTT).
- For copy-free transitions, perform these steps before initiating the Export & Halt 7-Mode Systems operation in the 7MTT.

## Steps

1. Open the ESXi host or vCenter Server managing the ESXi host.
2. Right-click the VM from which you need to remove snapshots.
3. Open the **Snapshot > Snapshot Manager** Snapshots window.
4. Click **Delete All**.

## Removing VM snapshots using the ESXi CLI

You can choose to use the ESXi CLI to remove the snapshots if you are using the Host Remediation Tool (HRT), or if you prefer the flexibility of CLI usage.

You must have the VMID from the Host VMs tab in the *Inventory Assessment Workbook* generated by the 7-Mode Transition Tool (7MTT).

## Steps

1. Use SSH to log in to the ESXi console.
2. Remove all VM snapshots of the VM with the applicable VMID:

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

After deleting snapshots, you must regenerate the *Inventory Assessment Workbook* to collect information related to Data ONTAP operating in 7-Mode and your ESXi hosts.

# Testing transitioned LUNs and ESXi host applications before the cutover phase of copy-based transitions

If you are using the 7-Mode Transition Tool (7MTT) 2.2 or later and Data ONTAP 8.3.2 or later to transition an ESXi host, you can test your transitioned clustered Data ONTAP LUNs to verify that you can bring your host and applications online before the cutover phase. Your source host can continue to run I/O to your source 7-Mode LUNs during testing.

- The new test host must be provisioned in a new test environment.

To avoid IP/MAC or UUID duplication or conflicts, test hosts must be configured in a private network.

- If you are transitioning a host booted from a local hard disk, the test host must have the same ESXi version and drivers as the source host.
- Zoning must be complete between FC or FCoE hosts and new clustered Data ONTAP nodes.
- Zoning must not exist between the source host and the new clustered Data ONTAP host.

If transitioned clustered Data ONTAP LUNs are visible to the source host during test mode, you might experience unexpected service disruptions on the source host.

- If you are transitioning a SAN booted host, your network adapter must be disabled.

You must maintain hardware parity between the test host and the source host and you must perform the following steps on the test host:



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

### Steps

1. After the baseline data copy is complete, select **Test Mode** in the 7MTT user interface (UI).
2. In the 7MTT UI, click **Apply Configuration**.
3. Zone the clustered Data ONTAP nodes on the test host.
4. Log in to the clustered Data ONTAP node, and then add new test host initiators to the igroup created by the 7MTT during the test phase.
5. Navigate to C:\Program Files\NetApp\operating in 7-Mode Transition Tool\.
6. Generate the 7-Mode to clustered Data ONTAP LUN mapping file from the Linux host where the 7MTT is installed:

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

For example:

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

7. Bring the test host online.
  - [Reconfigure your VMware software iSCSI initiator after transition](#)
  - [Set up your ESXi hosts configured for SAN boot after transition](#)
8. Verify that all transitioned clustered Data ONTAP LUNs are discovered.
9. If you are transitioning a non-SAN booted host, reregister your VMs.

[Re-registering VMs after transition of ESXi host.](#)

10. Complete the necessary post-transition steps for ESXi hosts.

[Post transition requirements for ESXi hosts](#)

11. Bring the host and applications online.
12. Perform your testing as needed.
13. Shut down the test host.
14. In the 7MTT UI, click **Finish Testing**.

The new clustered Data ONTAP LUNs are now read-only, and the data from the source 7-Mode LUNs is resynchronized.

15. If you plan to use the same source after you complete your transition, then edit the igroup on the clustered Data ONTAP node to add the appropriate initiator.

If you are planning to promote your test host to production, then you do not need to edit the igroup.

After you complete your test, plan time to shut down the source host connected to your controllers running Data ONTAP operating in 7-Mode. When you click **Complete Transition**, your source 7-Mode volumes and LUNs go offline and the transitioned clustered Data ONTAP LUNs become read/write.

## Related information

[Copy-based transition](#)

# Downtime at apply configuration (precutover) phase of ESXi host transition

You must plan for downtime at the apply configuration (precutover) phase while transitioning ESXi hosts.

After you complete the prerequisites for ESXi host transitions, you can use the 7-Mode Transition Tool (7MTT) to transition your LUNs from Data ONTAP operating in 7-Mode to ONTAP. ESXi hosts, VMs, and applications can be online until the apply configuration (precutover) phase of the 7MTT transition. However, at the apply configuration (precutover) phase, all applications and guest operating systems must be powered down. If ESXi hosts are SAN booted, or if SAN booted LUNs are part of the transition, the ESXi hosts must also be powered down at the apply configuration (precutover) phase.

If the host is not SAN booted and you need to keep your service running to non-transitioning LUNs or array LUNs that are not made by NetApp, you can choose not to shut it down. However, if you do not shut it down, you might experience the all-paths-down (APD) condition. To avoid the APD condition, you can mask Data ONTAP operating in 7-Mode LUNs. See the [VMware Knowledgebase ID 1009449](#) for more information.

## Post-transition remediation requirements for ESXi hosts

After you transition LUNs for ESXi hosts from Data ONTAP operating in 7-Mode to clustered Data ONTAP using the 7-Mode Transition Tool (7MTT), you must perform a series of remediation tasks to bring your LUNs online and begin servicing data.

## Related information

[Preparing for post-transition ESXi host remediation](#)

[Reregistering VMs after transition on non-SAN boot ESXi hosts using vSphere Client](#)

[Setting up ESXi hosts configured for SAN boot after transition](#)

[Determining whether VMFS volumes need to be remounted after transition](#)

[Reattaching RDM LUNs to VMs](#)

[Enabling CAW on a datastore using the ESXi CLI](#)

[Post-transition remediation for Linux and Windows guest operating systems](#)

[Recommended settings for ESXi hosts after transition remediation](#)

## Preparing for post-transition ESXi host remediation

After the 7-Mode Transition Tool (7MTT) transition is complete, you must perform various ESXi host remediation tasks. There are several steps you must complete before you can perform those tasks.

- For copy-based transitions (CBTs), perform these steps before initiating the Storage Cutover operation in the 7MTT.
- For copy-free transitions (CFTs), perform these steps before initiating the Export & Halt 7-Mode Systems operation in the 7MTT.

### Steps

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

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

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

For example:

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

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

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

For example:

```
transition cft export lunmap -p SanWorkLoad -s svm1 -o c:/Libraires/Documents/7-to-C-LUN-MAPPING-svm1.csv
```



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

2. Verify that igroup and initiator mappings are present.

7MTT re-creates the same igroup with initiators used in Data ONTAP operating in 7-Mode, and remaps the clustered Data ONTAP LUN to the host.

3. Verify that zoning is appropriate for the new clustered Data ONTAP target.

4. If you are doing a copy-free transition (CFT), run `vol rehost`.

See the [7-Mode Transition Tool Copy-Free Transition Guide](#) for `vol rehost` procedures.

## Reregistering VMs after transition on non-SAN boot ESXi hosts using vSphere Client

After you transition a non-SAN booted host, you must reregister your Virtual Machines (VMs).

The host must be online and the LUNs must be discovered.

- For copy-based transitions, perform these steps before initiating the Storage Cutover operation in the 7-Mode Transition Tool.
- For copy-free transitions, perform these steps before initiating the Export & Halt 7-Mode Systems operation in the 7MTT.

### Steps

1. Open the *Inventory Assessment Workbook* generated by the Inventory Collect Tool (ICT).

2. Navigate to the Host VMs tab, and then record the **VM Config File** path and **Location/Datastore Name** of the VM.
3. Use vSphere Client to log in to the ESXi host or the vCenter Server managing the ESXi host.
4. Under **Host and Clusters**, select the ESXi host.
5. Navigate to **Configuration > Hardware > Storage**.
6. Select the datastore with the datastore name you previously noted.
7. Right click and select **Browse Datastore**.

The Datastore Browser window opens.

8. Navigate to the **VM Config File** path you previously noted.
9. Right click the `.vmx` file, and then select **Add To Inventory**.
10. Repeat these steps for each VM listed in the **Host VMs** tab in the *Inventory Assessment Workbook* generated by the ICT.

## Reconfiguration of VMware software iSCSI initiator after transition

If your ESXi host accessed your Data ONTAP operating in 7-Mode system with the VMware software iSCSI initiator, then after transition from 7-Mode to clustered Data ONTAP, you must reconfigure the VMware software iSCSI initiator on your ESXi host and enable it to discover the new clustered Data ONTAP target.

For copy-based transitions, you must reconfigure your VMware software iSCSI initiator before initiating the Storage Cutover operation in the 7-Mode Transition Tool (7MTT). For copy-free transitions, you must reconfigure your VMware software iSCSI initiator before initiating the Export & Halt 7-Mode Systems operation in the 7MTT.

During the reconfiguration, you must retrieve the iSCSI IP and IQN used by the new clustered Data ONTAP target. If the target IP subnet has changed, the respective IP subnet changes are required on the host iSCSI initiator ports as well.

To make the required changes to the software iSCSI initiator on the VMware ESXi host, see the *VMware vSphere ESXi5.x Storage Guide*.

### Related information

[System administration](#)

## Setting up ESXi hosts configured for SAN boot after transition

If your ESXi host was configured for SAN boot before transition from Data ONTAP operating in 7-Mode, you must perform several steps before using the host after transition.

- For copy-based transitions, perform these steps before initiating the Storage Cutover operation in the 7MTT.
- For copy-free transitions, perform these steps before initiating the Export & Halt 7-Mode Systems operation in the 7MTT.

## Steps

1. Reconfigure your FC and FCoE HBA BIOS to boot from the SAN boot LUN of the clustered Data ONTAP system.
2. Boot the ESXi host.
3. Reset the host configurations to the pretransition settings.
4. For iSCSI hosts, see how to reconfigure your VMware iSCSI initiator.

[Reconfigure your VMware iSCSI initiator](#)

5. Remount the VMFS datastores created from the boot LUN in the default installation.

## Related information

[Remounting VMFS volumes after transition using vSphere Client](#)

[Remounting VMFS volumes after transition using the ESXi CLI](#)

[SAN administration](#)

## Determining whether VMFS volumes need to be remounted after transition

After transitioning from Data ONTAP operating in 7-Mode to clustered Data ONTAP, you might have VMFS volumes that need to be remounted to bring VMFS datastores and VMs to their pretransition states.

- For copy-based transitions, perform these steps before initiating the Storage Cutover operation in the 7-Mode Transition Tool (7MTT).
- For copy-free transitions, perform these steps before initiating the Export & Halt 7-Mode Systems operation in the 7MTT.

## Steps

1. Open the *Inventory Assessment Workbook* generated by the Inventory Collect Tool (ICT).
2. Click the **SAN Host Filesystems** tab.
3. Check the **Drive / Mount / Datastore Name** column for the filesystems and datastores mounted on the host before transition.
4. Make a note of the corresponding LUN naa IDs in the **SCSI Device ID/Device Name** column for the datastore.
5. Check whether the naa IDs noted for the datastore are listed in the 7MTT mapping file that was generated after the transition.
  - If none of the naa IDs are present in the 7MTT mapping file, the datastore and its underlying LUNs were not part of the 7MTT transition and no remediation is required.
  - If only part of the naa IDs are present in the 7MTT mapping file, your transition is incomplete and you cannot proceed.
  - If all naa IDs are present, you must remount your VMFS volumes.

## Related information

[Remounting VMFS volumes after transition using vSphere Client](#)

#### Remounting VMFS volumes after transition using vSphere Client

After transition, you must remount your VMFS volumes to bring your datastores and virtual machines (VMs) to their pretransition states. If you are not familiar with the ESXi CLI or it is more convenient in your environment, you can use vSphere Client to remount your volumes.

These steps apply to volumes and spanned volumes.

##### Steps

1. Log in to the ESXi host or the vCenter Server managing the ESXi host.
2. Under **Hosts and Clusters**, select the ESXi host.
3. Navigate to **Configuration > Hardware > Storage**.
4. In the upper-right corner, click **Add storage**.
5. Select **Disk/LUN**.
6. Click **Next**.
7. In the list of LUNs, locate the **VMFS\_label** column displaying the name of the datastore.
8. Select the LUN to complete the remount operation.

If you are remounting a spanned VMFS volume, the first LUN in the span is marked “head”. You must select the “head” LUN to complete the remount operation.

9. Click **Next**.
10. In the Select VMFS Mount Options window, select **Keep the existing signature**.
11. Complete the wizard.
12. Repeat these steps for all of the LUNs displaying the datastore name in the VMFS\_label column.

Datastores are remounted and VMs are active.

#### Remounting VMFS volumes after transition using the ESXi CLI

After transition, you can use the ESXi CLI to remount your volumes and bring your datastores and VMs to their pretransition states.

The original 7-Mode LUN must be unmapped or offline.

These steps apply to volumes and spanned volumes.

##### Steps

1. Log in to the ESXi console using SSH.
2. List the newly added LUNs with existing VMFS signatures and VMFS labels:

```
# esxcfg-volume -l
```

The following is an example of the LUNs listed with VMFS signatures and VMFS labels.

```
# 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. Remount the VMFS volumes persistently with the same signature:

- For regular volumes:  
**esxcfg-volume -M|--persistent-mount VMFS UUID|label**
- For spanned volumes:  
**# esxcfg-volume -M vmfs-span-ds**

## Related information

[VMware KB: vSphere handling of LUNs detected as snapshot LUNs](#)

## Reattaching RDM LUNs to VMs

For VMs attached to Raw Device Mapped (RDM) LUNs to function after transition, you must remove the RDM disks hosting the LUN from the VM. Then you must reattach the RDM disks to the VM based on the LUN serial number provided by the 7-Mode Transition Tool (7MTT).

- For copy-based transitions, perform these steps before initiating the Storage Cutover operation in the 7MTT.
- For copy-free transitions, perform these steps before initiating the Export & Halt 7-Mode Systems operation in the 7MTT.

## Steps

1. In the *Inventory Assessment Workbook*, navigate to the **Host VM Disk Details** tab.
2. Identify the ESXi host VM with PTRDM or NPTRDM in the **Type** column.
3. Note the VM name, the disk path details in the **Disk** column, and the naa ID from the **Device Mapped** column.
4. Verify that the naa ID is listed in the 7MTT mapping file that was generated after transition.
5. Verify that the naa ID has a corresponding new naa ID in the **LUN WWID** column of the mapping file.

This is the new clustered Data ONTAP LUN naa ID.

6. Use the clustered Data ONTAP LUN naa ID from the **LUN WWID** column and the disk path details to reattach the clustered Data ONTAP LUN to the VM.

## Related information

[Removing stale RDMs using vSphere Client](#)

## Removing stale RDMs using vSphere Client

All RDM LUNs become stale during the transition from ONTAP operating in 7-Mode to clustered Data ONTAP. After transition, the RDMs must be removed and reattached before the LUNs can begin servicing data.

You must have the VM name and disk path of the RDM from the *Inventory Assessment Workbook*.

### Steps

1. Open the ESXi host or vCenter Server managing the ESXi host.
2. Right-click the VM, and then select **Edit Settings**.

The VM Properties window is displayed.


3. Select the hard disk from the list of devices using the disk path from the *Inventory Assessment Workbook*.
4. Make note of the **Virtual Device Node** and **Compatibility Mode** from the VM Properties window.

Virtual Device Node: SCSI 0:2

Compatibility Mode: Physical

5. Click **Remove**.
6. Select **Remove from virtual machine and delete files from disk**.
7. Click **OK**.

An error message similar to the following appears. You can ignore this message.



8. Click **Close**.

## Reattaching RDM to VMs using vSphere Client

After transition of an ESXi host using the 7-Mode Transition Tool (7MTT), you must reattach your RDMs to virtual machines (VMs).

Your stale Raw Device Mappings (RDMs) must have been removed.



## Steps

1. Open the ESXi host or vCenter server managing the ESXi host.
2. Right-click the VM, and then select **Edit Settings**.

The VM Properties window opens.

3. Click **Add**.

The Add Hardware window opens.

4. Click **Hard Disk**.
5. Click **Next** to select the disk.
6. Select **Raw Device Mappings**.
7. Click **Next** to select the target LUN.
8. Select the LUN with the new clustered Data ONTAP NAA ID that you noted from the 7MTT mapping file.
9. Click **Next**.
10. Choose **Select Datastore**.
11. Select the datastore that matches the disk path that you noted in the 7MTT mapping file.
12. Click **Next**.
13. Choose either **Physical** or **Virtual** for the **Compatibility Mode**.

Choose the compatibility mode noted when your stale RDM was removed.

14. Click **Next**.
15. Choose **Advance Options**.
16. Select the **Virtual Device Node**.

Select the virtual device mode that you noted when you removed the stale RDM.

17. Click **Next**.
18. Click **Finish** to submit your changes.
19. Repeat the steps for all VMs with RDM attached.

## Related information

[Removing stale RDMs using vSphere Client](#)

## Reattaching RDM using ESXi the CLI/console

After transition from Data ONTAP operating in 7-Mode to clustered Data ONTAP you must reattach your raw device mapping (RDM).

- You must retrieve the RDM disk file listed in the Disk column of the *Inventory Assessment Workbook*.
- You must retrieve the new clustered Data ONTAP LUN naa ID from the 7MTT mapping file.

## Steps

1. Log in to the ESXi console using SSH.

2. Use the mv command to make a backup of the RDM disk file and the associated device file.

The RDM disk file is listed in the Disk column of the *Inventory Assessment Workbook*.

If the RDM disk file is /vmfs/volumes/53a3ac3d-df5aca03-3a94-001ec9d631cb/VM2-win-bus-A/VM2-winbus-A.vmdk, you would issue the following command:

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

- For physical compatibility RDM (PTRDM):

```
mv RDM_disk_file_name-rdmp.vdmk RDM_disk_file_name-rdmp.vdmk_bak
```

For example:

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

- For virtual compatibility RDM (NPTRDM):

```
mv RDM_disk_file_name-rdmp.vdmk RDM_disk_file_name-rdmp.vdmk_bak
```

For example:

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

3. Use the new clustered Data ONTAP LUN naa ID and the RDM disk file to re-create the RDM configuration and device files.

- For PTRDM:

```
# vmkfstools -z /vmfs/devices/disks/new_clustered_Data_ONTAP_naa_ID.vmdk
```

For example:

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

- For NPTRDM:

```
# vmkfstools -r /vmfs/devices/disks/new_clustered_Data_ONTAP_naa_ID.vmdk
```

For example:

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

4. 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 are displayed under the VM directory path.

5. Repeat the steps for all VMs with RDM attached.
6. Restart the hostd and vpxa agents in the ESXi host:

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

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

## Post-transition remediation for Linux and Windows guest operating systems

Linux and Windows guest operating systems might require additional remediation after transition of LUNs from Data ONTAP operating in 7-Mode to clustered Data ONTAP.

For copy-based transitions, do the following after completing the Storage Cutover operation in the 7MTT. For copy-free transitions, do the following after the Import Data & Configuration operation in the 7MTT is complete.

- Linux

If the mount points are defined in the `/etc/fstab` file, you must mount the LUN (`mount --a`).

- Windows

If Failover Cluster is configured on the VM, you must bring the disks online from Failover Cluster Manager.

## Recommended settings for ESXi hosts after transition remediation

After you complete the post-transition remediation steps for your ESXi host, you should apply the recommended ESXi host settings for clustered Data ONTAP on the host.

You can use Virtual Storage Console (VSC) to configure ESXi host settings. VSC is the standard NetApp plugin that enables vSphere vCenter to configure ESXi host settings for Data ONTAP. ESXi hosts and virtual machines (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 Knowledgebase 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*

## Enabling CAW on a datastore using the ESXi CLI

If you did not have support for Compare and Write (CAW) in Data ONTAP operating in 7-Mode, you must manually enable CAW support when you transition to clustered Data ONTAP. Clustered Data ONTAP supports CAW by default.

- There must be no I/O or VMs actively running on the VMFS datastore.
- The datastore must be remounted, if it was migrated.
- You must have the new ONTAP LUN naa ID from the 7-Mode Transition Tool (7MTT) mapping file.

CAW must be enabled only when no I/O or VMs are actively running on the VMFS datastore.

- CAW must be enabled only when no I/O or VMs are actively running on the VMFS datastore.
- For copy-based transitions, perform these steps after completing the Storage cutover operation in the 7MTT.
- For copy-free transitions, perform these steps after the Import Data & Configuration operation in the 7MTT is complete.

## Steps

1. Open the *Inventory Assessment Workbook* generated by the Inventory Collect Tool (ICT).
2. Navigate to the SAN Host Filesystems tab.
3. Verify the CAW status for the datastore.

The **ATS/CAW** value for the datastore should display **Disabled**, and the Filesystem column should display **VMFS.x**.

4. Note the name of the datastore in the Disk column
5. Log in to the ESXi console using SSH.
6. List the device and partition details:

```
~ # vmkfstools -Ph -v1 datastore_path
```

*datastore\_path* is the datastore name from the Disk column of the *Inventory Assessment Workbook*.

```
# vmkfstools -Ph -v1 /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

```

7. Note the first device name and partition number.

In the preceding example, naa.600a098044314c6c442b446d51376749:1 is the device name and partition number.

8. Use the device ID and partition number to enable CAW on the datastore:

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

9. Verify that the VMFS volume has been configured with ATS only:

```
# vmkfstools -Ph -v1 /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

```

## Related information

[Remounting VMFS volumes after transition using vSphere Client](#)

[Remounting VMFS volumes after transition using the ESXi CLI](#)

[VMware documentation](#)

## Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

## Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.