



Post-transition remediation requirements for ESXi hosts

ONTAP 7-Mode Transition

NetApp
February 11, 2024

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

Table of Contents

- Post-transition remediation requirements for ESXi hosts. 1
 - Preparing for post-transition ESXi host remediation 1
 - Reregistering VMs after transition on non-SAN boot ESXi hosts using vSphere Client 2
 - Reconfiguration of VMware software iSCSI initiator after transition. 3
 - Setting up ESXi hosts configured for SAN boot after transition 3
 - Determining whether VMFS volumes need to be remounted after transition 4
 - Reattaching RDM LUNs to VMs. 6
- Post-transition remediation for Linux and Windows guest operating systems 10
- Recommended settings for ESXi hosts after transition remediation 10

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

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

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](#)

[Reattaching RDM to VMs using vSphere Client](#)

[Reattaching RDM using ESXi the CLI/console](#)

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-win-bus-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-win-bus-  
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-win-bus-  
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*

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.