



# Preparing for transition of ESXi hosts

## ONTAP 7-Mode Transition

NetApp  
May 31, 2021

# Table of Contents

- Preparing for transition of ESXi hosts ..... 1
  - What the Inventory Collect Tool is ..... 1
  - Preparing Linux guest operating systems for transition ..... 2
  - Preparing Windows guest operating systems for transition ..... 4
  - How to identify VM snapshots that must be removed before transition ..... 5

# 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.

## Copyright Information

Copyright © 2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system- without prior written permission of the copyright owner.

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

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

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

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

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## Trademark Information

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