



HPE Morpheus VM Essentials Software Migration Guide

Step-by-step technical guide for users migrating VMs from VMware to HPE HVM. Includes Linux (Ubuntu, Red Hat Enterprise Linux) and Microsoft Windows Operating System.

Contents

Introduction.....	3
Prerequisites and requirements	3
Migration plans	4
Migrating the VM	4
Ubuntu	7
Supported versions	7
Prepare the source VM.....	7
Migrating the VM on Ubuntu.....	10
Post migration	11
Rollback plan	15
Red Hat Enterprise Linux 8.x	15
Supported versions	15
Prepare the source VM.....	15
Migrating the VM on Red Hat.....	19
Post migration Steps	19
Rollback plan	22
Windows Server 2022 (or later).....	22
Supported versions and assumptions	22
Prepare the OS	22
Migrating the VM on Windows.....	27
Rollback plan	27
Troubleshooting	27
Conclusion.....	28
Appendix	29
Procedure to create an external datastore for images and VMs.....	29
Reference links	31

Introduction

HPE Morpheus VM Essentials Software is a new virtualization software solution that allows customers to provision and manage HPE Morpheus VM Essentials hypervisor VMs (KVM) and VMware®-based virtual machines (VMs) from a single intuitive interface. The software natively provides tooling to migrate VMware VMs to HVM Clusters. This feature supports unidirectional migration of VMs from VMware to the HPE Morpheus VM Essentials Software hypervisor.

This document outlines in detail the offline migration procedure for Ubuntu 22.04 (or later), Red Hat® Enterprise Linux® (RHEL) 8.x (or later), and Windows 2022 (or later) OS-based VMware virtual machines to HPE Morpheus VM Essentials Software.

The information in this guide is intended for virtualization administrators, as a step-by-step guide to help them migrate VMs to HPE Morpheus VM Essentials Software. This document assumes that readers are experienced in performing common OS, storage, and network administration tasks.

Prerequisites and requirements

HPE Morpheus VM Essentials Software v8.0.8-x is installed and configured with items listed in the following:

- HVM Hosts must be upgraded to Agent version 2.10.0 at minimum. Users who installed HPE Morpheus VM Essentials on a lower version should check and, if necessary, upgrade.
- HVM Hosts must be able to reach VMware ESXi™ Hosts and VMware vCenter® on the target VMware vCenter Cloud via the Management Network.
- The Source VMs must be running for the preparation phase of the migration to complete successfully. If VMs are not running, they will automatically be restarted.

First release requirements:

This section includes requirements for the current version of the migration feature which are subject to change and improvement as road-mapped enhancements are included with subsequent versions of the product.

- The destination datastore must have enough “thick” space for each VM.
- VMs will power down prior to the transfer, meaning service of the source workload will be disrupted during the transfer.
- VirtIO / VirtIO-SCSI target storage must be supported by the source VM.
- Currently supported operating systems: Red Hat, CentOS, Rocky, Alma, SUSE, Ubuntu, Debian, Windows (currently requires manual preparation steps described below).
- Source VMs must be capable of getting the qemu-guest-tools package installed via package managers (yum, apt, dnf, zypper). If they are not, this can be manually performed by setting the "skip guest tools" flag on the migration.
- RDM (Raw Device Mappings) are not yet supported.
- Source VMs must not have any attached ISOs or cdroms. These are not supported and the migration will fail.

Recommendations:

- Batch limits and bandwidth limitations testing is still in progress. It's currently recommended you migrate no more than 20 VMs at a time. Testing is still ongoing to determine the upper limits of migrations and this recommendation is likely to increase over time.
- Begin using this feature with a smaller migration than the limit to make sure your workloads are moving correctly.

Note

Before migrating any VMs, it is recommended to back up the source VM and have a rollback plan in place.








Migration plans

Migration plans are created and run from the Migrations section of the Tools menu (**Tools > Migrations**). Migrations happen by creating and running Migration Plans. Plans are created in a pending state, meaning an additional action must be undertaken to set the plan in motion. Thus, Plans may be created which are intended to be run later. Once finished, completed Plans remain on the Migrations list page for later review and, if desired, deletion.

Migrating the VM

Begin creating a Migration Plan from the Migrations list page (**Tools > Migrations**). This page contains a list of all Migrations already created, including those completed, currently running, and available to be run (pending). To start a new Migration, click **+ ADD**.

MIGRATIONS

NAME	SOURCE	TARGET	RESOURCE POOL	VMS	STATUS	
Alletra Migration	QA VMware	HPE Alletra VME	VME Alletra MP Cluster	0	COMPLETED	
auto-vmw2hvm-0721	QA VMware	QA MorpheusVM	QA MVM Cluster	0	COMPLETED	
Big Boy	QA VMware	QA MorpheusVM	QA MVM Cluster	0	RUNNING	
Davids Migration Plan	QA VMware	QA MorpheusVM	QA MVM Cluster	0	COMPLETED	
Migration Multi	QA VMware	QA MorpheusVM	QA MVM Cluster	1	COMPLETED	
Rock Test	QA VMware	QA MorpheusVM	QA MVM Cluster	0	COMPLETED	
TM Temp	QA VMware	QA MorpheusVM	QA MVM Cluster	1	PENDING	

From the **SETUP** tab of the **CREATE MIGRATION PLAN modal**, configure the following:

- **NAME:** A name to identify the Migration
- **SOURCE:** The source VMware vCenter Cloud
- **TARGET:** The target Cloud containing the HVM Cluster
- **RESOURCE POOL:** The selected destination HVM Cluster
- **GROUP:** The Group which should own the migrated VMs

When finished, click **NEXT**.

CREATE MIGRATION PLAN ✕

SETUP
>
CHOOSE VMS
>
MAP RESOURCES
>
REVIEW

Setup Migration Plan

NAME

SOURCE

TARGET

RESOURCE POOL

GROUP

PREVIOUS
NEXT

The next step is to choose VMs from the selected VMware source Cloud to migrate. Select as many as desired given current recommendations regarding the maximum size due to storage space and available bandwidth (see the previous section). Selected VMs will form a list at the bottom of the modal. When finished, click **NEXT**.

CREATE MIGRATION PLAN ✕

SETUP
>
CHOOSE VMS
>
MAP RESOURCES
>
REVIEW

SELECT VMS

SELECT PAGE
SELECT ALL

POWER	OS	VM NAME	
		aw-feat-vmw-vmware-1	SELECT
		aw-instance-98	SELECT
		aw-instance-99	SELECT
		BungeWorker	SELECT
		CentOS-VersionTest	SELECT

The next tab establishes resource mapping. Listed, you will see the current networks and storage locations for the selected VMs. Choose destination networks and destination storage locations for each. Optionally, enter existing Linux or Windows user credentials and choose if prechecks or guest tools installation should be skipped. Click **NEXT**.

CREATE MIGRATION PLAN
✕

SETUP
CHOOSE VMS
MAP RESOURCES
REVIEW

Networks

SOURCE	TARGET
VLAN0043 - QA	Compute ▼

Storage

SOURCE	TARGET
ESXi-DC2-QA-LUN01	mvm-qa-vol01 (4092.39 Gib) ▼

Linux Settings

USERNAME

PASSWORD

SSH KEY Select ▼

Migration Options

SKIP PRECHECKS

SKIP GUEST TOOLS

PREVIOUS
NEXT

The final tab is a review tab where all current selections can be checked. Return to any previous tabs, if necessary, to update selections.

At this point, the Migration is created but will not run without additional input. The Migration is in a "Pending" state. All selections made when configuring the Migration are shown here. To execute the Migration, click **RUN**. The length of time it will take for the Migration to run depends on many factors. Once run, the Migration detail page will provide status updates on VMs currently being migrated, which have completed successfully, and if any have failed. The History tab provides greater detail on current and past migration actions taken. The Destination tab includes details on VMware VMs which have successfully been migrated to the HVM Cluster. Migrations are run just once and may be kept indefinitely after for review. When a Migration is no longer needed, click **DELETE**.

Migrations > TM Temp

RUN
DELETE

TM Temp

Status: PENDING Source Cloud: QA VMware Destination Cloud: QA MorpheusVM Destination Resource Pool: QA MVM Cluster Group: QA MorpheusVM

1

PENDING

0

IN PROGRESS

0

FAILED

0

READY

1

TOTAL

Progress

Source

Destination

History

SERVERS

OS	SOURCE	DESTINATION	STATUS
qa-feat-cypher		Unknown	PENDING

NETWORKS

SOURCE	DESTINATION
VLAN0002 - Internal Server	Compute VLAN 43

DATASTORES

SOURCE	DESTINATION
ESXi-DC2-QA-LUN02	mvm-gfs2

Ubuntu

Supported versions

– 22.04 (or later)

Prepare the source VM

Important

The HPE Morpheus VM Essentials bulk migration tool is designed to handle Ubuntu-based workloads automatically in most cases. The preparation steps listed here can be taken as a troubleshooting measure if bulk migrations of Ubuntu-based workloads are failing. In a typical scenario, you should be able to skip over this section on preparing the source VM and go straight to the next section on migrating VMs.

Before the start of the migration, there are prerequisites to complete on the source virtual machine:

- Document source VM storage devices and mount points.
- Record network configuration details.
- Verify if VirtIO drivers and qemu-guest-agent are installed.
- Eject CD-ROM/DVD-ROM.

The following section describes the preparation steps for the source VM before migration.

1. Record the filesystem mount points and other disk information by running the following commands:

```
cat /etc/fstab
lsblk
df -h
ls -lah /dev/disk/by-uuid
```

- a. **Recommended**

The mounts can be changed from using the device name (ex. /dev/sdb) to using their ID, UUID, or other static values if they are not already. This will help ensure the mounts work without intervention after the migration. For example, using UUID:

```
ls -lah /dev/disk/by-uuid
sse@ubuntuvmetest:~$ ls -lah /dev/disk/by-uuid/

total 0

drwxr-xr-x 2 root root 140 Mar 18 19:35 .
drwxr-xr-x 7 root root 140 Mar 10 06:59 ..
lrwxrwxrwx 1 root root 10 Mar 10 17:50 608e457f-a9c0-4594-ad37-6fcbe699e601 ->
../../dm-0
lrwxrwxrwx 1 root root 10 Mar 10 17:50 99EA-150F -> ../../sda1
lrwxrwxrwx 1 root root 10 Mar 10 17:50 b228a623-3515-41a8-a7e2-acf6ab86e099 ->
../../sda2
lrwxrwxrwx 1 root root 10 Mar 18 19:35 e85a554d-78a9-4e9f-a15d-d8ba4b6031bd ->
../../sdc1
lrwxrwxrwx 1 root root 10 Mar 18 19:35 f5346b5c-becf-400c-affe-2730ce19ced4 ->
../../sdb1

/etc/fstab
sse@ubuntuvmetest:~$ cat /etc/fstab

# /etc/fstab: static file system information.
```

```

#
# Use 'blkid' to print the universally unique identifier for a
# device; this may be used with UUID= as a more robust way to name devices
# that works even if disks are added and removed. See fstab(5).
#
# <file system> <mount point> <type> <options> <dump> <pass>
# / was on /dev/ubuntu-vg/ubuntu-lv during curtin installation
/dev/disk/by-id/dm-uuid-LVM-
7oPWnZJPEPJxmCpr2eGLKPdhE7fJZdk9UEkhdv4Uvg58SMYCHW4FIRy18SrujYCD / ext4 defaults 0 1
# /boot was on /dev/sda2 during curtin installation
/dev/disk/by-uuid/b228a623-3515-41a8-a7e2-acf6ab86e099 /boot ext4 defaults 0 1
# /boot/efi was on /dev/sda1 during curtin installation
/dev/disk/by-uuid/99EA-150F /boot/efi vfat defaults 0 1
/swap.img none swap sw 0 0
/dev/disk/by-id/f5346b5c-becf-400c-affe-2730ce19ced4 /disk1 xfs defaults 0 1
/dev/disk/by-id/e85a554d-78a9-4e9f-a15d-d8ba4b6031bd /disk2 xfs defaults 0 1

```

b. Note if additional disks are mounted using the device name, such as:

```

/dev/sdb
/dev/sdc

```

c. If the devices remain mounted using the device names, **the mounts will need to be fixed after migration**, as the new devices will be **/dev/vdb**, **/dev/vdc**, and such. Additional details are listed in the ["Post migration"](#) section.

For example:

```
cat /etc/fstab
```

```
# /etc/fstab: static file system information.
```

```
#
```

```

# Use 'blkid' to print the universally unique identifier for a
# device; this may be used with UUID= as a more robust way to name devices
# that works even if disks are added and removed. See fstab(5).
#

```

```
#
```

```
# <file system> <mount point> <type> <options> <dump> <pass>
```

```
# / was on /dev/ubuntu-vg/ubuntu-lv during curtin installation
```

```

/dev/disk/by-id/dm-uuid-LVM-
7oPWnZJPEPJxmCpr2eGLKPdhE7fJZdk9UEkhdv4Uvg58SMYCHW4FIRy18SrujYCD / ext4 defaults 0 1

```

```
# /boot was on /dev/sda2 during curtin installation
```

```
/dev/disk/by-uuid/b228a623-3515-41a8-a7e2-acf6ab86e099 /boot ext4 defaults 0 1
```

```
# /boot/efi was on /dev/sda1 during curtin installation
```

```
/dev/disk/by-uuid/99EA-150F /boot/efi vfat defaults 0 1
```

```
/swap.img none swap sw 0 0
```

```

/dev/sdb1      /disk1  xfs      defaults 0      1
/dev/sdc1      /disk2  xfs      defaults 0      1

```

2. Note/save the networking connection device names and settings:

```

ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
inet 127.0.0.1/8 scope host lo
valid_lft forever preferred_lft forever
inet6 ::1/128 scope host
valid_lft forever preferred_lft forever

2: ens160: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
link/ether 00:50:56:9d:21:05 brd ff:ff:ff:ff:ff:ff
inet 10.3.23.223/24 metric 100 brd 10.3.23.255 scope global dynamic ens160
valid_lft 4300sec preferred_lft 4300sec
inet6 fe80::250:56ff:fe9d:2105/64 scope link
valid_lft forever preferred_lft forever

```

```

cat /etc/netplan/*
# This is the network config written by 'subiquity'
network:
  ethernets:
    ens160:
      dhcp4: true
  version: 2

```

- a. After the migration, the networking device names may change and will need to be updated in any **/etc/netplan/** files.
- b. Once the netplan files have been updated, a **netplan apply** will need to be run.
- c. Example before migration:

```

# This is the network config written by 'subiquity'
network:
  ethernets:
    ens160:
      dhcp4: true
  version: 2

```

- d. Example after migration:

```

# This is the network config written by 'subiquity'
network:
  ethernets:

```

```
enp3s0:
```

```
dhcp4: true
```

```
version: 2
```

e. Additional details are listed in the [“Post migration”](#) section.

3. Verify if VirtIO drivers are already installed on the operating systems:

```
grep -e "^CONFIG_VIRTIO_NET=.*" -e "^CONFIG_VIRTIO_BLK=.*" -e "^CONFIG_VIRTIO_PCI=.*"
/boot/config-$(uname -r)
CONFIG_VIRTIO_BLK=y

CONFIG_VIRTIO_NET=y

CONFIG_VIRTIO_PCI=y
```

a. All values should return with y.

b. If another value was seen other than y, check if the initrd.img temporary file systems contains the drivers needed:

```
lsinitramfs /boot/initrd.img-$(uname -r) | grep -e virtio_blk.ko -e virtio_scsi.ko
usr/lib/modules/5.15.0-134-generic/kernel/drivers/block/virtio_blk.ko

usr/lib/modules/5.15.0-134-generic/kernel/drivers/scsi/virtio_scsi.ko
```

l. If the drivers virtio_blk.ko and virtio_scsi.ko are seen as highlighted above, the migration should complete successfully and the VM should boot successfully.

c. If the drivers cannot be successfully detected in initrd.img, the VirtIO drivers may need to be added to the operating system and the initrd.img temporary file system:

l. These are example commands that may help with including the drivers in initrd.img, but consulting the OS documentation is recommended:

```
echo -e "virtio_blk\nvirtio_pci\nvirtio_blk_net" | tee -a /etc/initramfs-
tools/modules

update-initramfs -u
```

4. If not already installed, install the qemu-guest-agent package:

```
apt install qemu-guest-agent -y
```

5. If one does not already exist, it is recommended to create a local user, which can be used to access the new virtual machine. In cases where the VM uses authentication from a remote source such as Active Directory, it may be unavailable if the network is not available immediately.

6. Eject any CD-ROM/DVD-ROMs that may be attached through VMware by selecting **Client Device** from the interface.




Figure 1. Eject CD

7. Shut down the source VM to ensure data consistency.

Migrating the VM on Ubuntu

Refer back to the section earlier in this document on creating [Migration Plans](#). The HPE Morpheus VM Essentials Bulk Migration tool can easily and automatically migrate VMs to the HVM hypervisor. In most cases, Ubuntu-based workloads do not require any preparation steps.

Post migration

After migration, there may be additional items that need to be resolved, though in most cases the bulk migration tool can handle these post-migration steps. If desired, you may check the following during your initial test migrations to have an idea of what you may need to do following migrations of production workloads. The following are a few common items seen (some actions would require a privileged user account):

1. Network adapter name updates:

- a. As mentioned in the [“Prepare the OS”](#) section, the networking device names may change and will need to be updated in any `/etc/netplan/` files.
- b. Once the netplan files have been updated, a **netplan apply** will need to be run.
- c. Example before migration:

```
cat /etc/netplan/*
# This is the network config written by 'subiquity'

network:

  ethernets:

    ens160:

      dhcp4: true

  version: 2
```

d. Example after migration:

```
ip a

1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever

2: enp3s0: <BROADCAST,MULTICAST> mtu 1500 qdisc mq state DOWN group default qlen 1000
    link/ether 00:50:56:9d:21:05 brd ff:ff:ff:ff:ff:ff

cat /etc/netplan/*
# This is the network config written by 'subiquity'

network:

  ethernets:

    enp3s0:

      dhcp4: true

  version: 2
```

e. Run **netplan apply**.

f. Result:

```
ip a
```

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default
qlen 1000

link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00

inet 127.0.0.1/8 scope host lo

valid_lft forever preferred_lft forever

inet6 ::1/128 scope host

valid_lft forever preferred_lft forever

2: enps3s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group
default qlen 1000

link/ether 00:50:56:9d:21:05 brd ff:ff:ff:ff:ff:ff

inet 10.3.23.223/24 metric 100 brd 10.3.23.255 scope global dynamic ens160

valid_lft 4300sec preferred_lft 4300sec

inet6 fe80::250:56ff:fe9d:2105/64 scope link

valid_lft forever preferred_lft forever
```

2. Device mounting when using device names that are not UUID or other static values:

a. This information is used if disks are mounted using the device name, such as:

```
/dev/sdb
/dev/sdc
```

If the devices remain mounted using the device names, **the mounts will need to be fixed after migration**, as the new devices will be **/dev/vdb, /dev/vdc**, and such. For example:

```
cat /etc/fstab
```

```
# /etc/fstab: static file system information.
```

```
#
```

```
# Use 'blkid' to print the universally unique identifier for a
```

```
# device; this may be used with UUID= as a more robust way to name devices
```

```
# that works even if disks are added and removed. See fstab(5).
```

```
#
```

```
# <file system> <mount point> <type> <options> <dump> <pass>
```

```
# / was on /dev/ubuntu-vg/ubuntu-lv during curtin installation
```

```
/dev/disk/by-id/dm-uuid-LVM-
```

```
7oPWnZJPEPJxmCpr2eGLKpdhE7fJZdk9UEkhdv4Uvg58SMYCHW4FIRy18SrujYCD / ext4 defaults 0 1
```

```
# /boot was on /dev/sda2 during curtin installation
```

```
/dev/disk/by-uuid/b228a623-3515-41a8-a7e2-acf6ab86e099 /boot ext4 defaults 0 1
```

```
# /boot/efi was on /dev/sda1 during curtin installation
```

```
/dev/disk/by-uuid/99EA-150F /boot/efi vfat defaults 0 1
```

```

/swap.img      none    swap    sw      0      0
/dev/sdb1      /disk1  xfs     defaults 0      1
/dev/sdc1      /disk2  xfs     defaults 0      1

```

- b. The new device names should align closely to the same, but first verify the available disks:

```
sse@ubuntuvmtest:~$ lsblk
```

```

NAME                                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINTS
vda                                  8:0    0   16G  0 disk
├─vda1                               8:1    0   763M  0 part /boot/efi
├─vda2                               8:2    0    1.8G  0 part /boot
└─vda3                               8:3    0   13.5G  0 part
   └─ubuntu--vg-ubuntu--lv          253:0  0    10G  0 lvm  /
vdb                                  8:16   0    10G  0 disk
└─vdb1                              8:17   0    10G  0 part /disk1
vdc                                  8:32   0    11G  0 disk
└─vdc1                              8:33   0    11G  0 part /disk2

```

- c. Prior to migration, the device names were **sdb** and **sdc** but are now **vdb** and **vdc**, respectively. However, they are not the same; be sure to map them accordingly.
- l. Additionally, output saved from running **ls -lah /dev/disk/by-uuid** while going through the steps in the Prepare the source VM section can help identify the disks, as the UUID should not change. For example:

ls -lah /dev/disk/by-uuid

Premigration:

```

ls -lah /dev/disk/by-uuid
sse@ubuntuvmtest:~$ ls -lah /dev/disk/by-uuid/
total 0
drwxr-xr-x 2 root root 140 Mar 18 19:35 .
drwxr-xr-x 7 root root 140 Mar 10 06:59 ..
lrwxrwxrwx 1 root root 10 Mar 10 17:50 608e457f-a9c0-4594-ad37-6fcbe699e601 ->
../../dm-0
lrwxrwxrwx 1 root root 10 Mar 10 17:50 99EA-150F -> ../../sda1
lrwxrwxrwx 1 root root 10 Mar 10 17:50 b228a623-3515-41a8-a7e2-acf6ab86e099 ->
../../sda2
lrwxrwxrwx 1 root root 10 Mar 18 19:35 e85a554d-78a9-4e9f-a15d-d8ba4b6031bd ->
../../sdc1
lrwxrwxrwx 1 root root 10 Mar 18 19:35 f5346b5c-becf-400c-affe-2730ce19ced4 ->
../../sdb1

```

Post migration:

```

ls -lah /dev/disk/by-uuid
sse@ubuntuvmtest:~$ ls -lah /dev/disk/by-uuid/
total 0

```

```

drwxr-xr-x 2 root root 140 Mar 10 20:10 .
drwxr-xr-x 7 root root 140 Mar 10 20:10 ..
lrwxrwxrwx 1 root root 10 Mar 10 20:10 608e457f-a9c0-4594-ad37-6fcbe699e601 ->
../../dm-0
lrwxrwxrwx 1 root root 10 Mar 10 20:10 99EA-150F -> ../../vda1
lrwxrwxrwx 1 root root 10 Mar 10 20:10 b228a623-3515-41a8-a7e2-acf6ab86e099 ->
../../vda2
lrwxrwxrwx 1 root root 10 Mar 18 20:10 e85a554d-78a9-4e9f-a15d-d8ba4b6031bd ->
../../vdc1
lrwxrwxrwx 1 root root 10 Mar 18 20:10 f5346b5c-becf-400c-affe-2730ce19ced4 ->
../../vdb1

```

Once all device names are sorted, update the **/etc/fstab** with the appropriate device names and partition, for example:

```

cat /etc/fstab

# /etc/fstab: static file system information.
#
# Use 'blkid' to print the universally unique identifier for a
# device; this may be used with UUID= as a more robust way to name devices
# that works even if disks are added and removed. See fstab(5).
#
# <file system> <mount point> <type> <options> <dump> <pass>
# / was on /dev/ubuntu-vg/ubuntu-lv during curtin installation
/dev/disk/by-id/dm-uuid-LVM-
7oPwnZJPEPJxmCpr2eGLKPdhE7fJZdk9UEkhdv4Uvg58SMYCHW4FIRy18SrujYCD / ext4 defaults 0 1
# /boot was on /dev/sda2 during curtin installation
/dev/disk/by-uuid/b228a623-3515-41a8-a7e2-acf6ab86e099 /boot ext4 defaults 0 1
# /boot/efi was on /dev/sda1 during curtin installation
/dev/disk/by-uuid/99EA-150F /boot/efi vfat defaults 0 1
/swap.img none swap sw 0 0
/dev/vdb1 /disk1 xfs defaults 0 1
/dev/vdc1 /disk2 xfs defaults 0 1

```

d. Once the mounts are updated correctly, use the following to mount them:

```

mount -a
systemctl daemon-reload

```

e. Using `df -h` or other tools should now show the partitions mounted to the locations specified.

```

df -h

Filesystem                Size  Used Avail Use% Mounted on
tmpfs                    197M  1.3M  196M   1% /run
/dev/mapper/ubuntu--vg-ubuntu--lv 9.8G  5.0G  4.3G  54% /

```

```

tmpfs                982M    0  982M    0% /dev/shm
tmpfs                5.0M    0  5.0M    0% /run/lock
/dev/vda2            1.7G   127M   1.5G    8% /boot
/dev/vda1            762M    6.1M   756M    1% /boot/efi
tmpfs                197M    4.0K   197M    1% /run/user/1000
/dev/vdb1            10G    104M   9.9G    2% /disk1
/dev/vdc1            11G    111M   11G     1% /disk2

```

Rollback plan

If the VM fails to come up on HPE Morpheus VM Essentials Software, shut down the VM on the software hypervisor. Power on the source VM either from vCenter or the software user interface.

Red Hat Enterprise Linux 8.x

Supported versions

- RHEL 8.x or later

Prepare the source VM

Before the start of the migration, there are prerequisites to complete on the source virtual machine:

- Document source VM storage devices and mount points.
- Record network configuration details.
- Verify if VirtIO drivers and qemu-guest-agent are installed.
- Eject CD-ROM/DVD-ROM.

The following section describes the preparation of source VM before migration.

1. Record the filesystem mount points and other disk information by running the following commands:

```

cat /etc/fstab
lsblk
df -h
ls -lah /dev/disk/by-uuid

```

a. Recommended

The mounts can be changed from using the device name (ex. **/dev/sdb**) to using their ID, UUID, or other static values if they are not already. This will help ensure the mounts work without intervention after the migration. For example, using UUID:

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

```

total 0
drwxr-xr-x. 2 root root 180 Mar 19 06:45 .
drwxr-xr-x. 8 root root 160 Feb 28 11:02 ..
lrwxrwxrwx. 1 root root  10 Feb 28 11:02 1970-1C3F -> ../../sda1
lrwxrwxrwx. 1 root root   9 Feb 28 11:02 2022-06-28-07-51-01-00 -> ../../sr0
lrwxrwxrwx. 1 root root  10 Feb 28 11:02 44e3e59c-fc46-4c1a-808c-e2d31b8fc57b ->
../../dm-1
lrwxrwxrwx. 1 root root  10 Feb 28 11:02 768394c6-7eca-48ad-a711-c93588e1b4f6 ->
../../dm-0

```

```
lrwxrwxrwx. 1 root root 10 Mar 19 06:45 c5fc2403-d7a2-40db-8046-76ef3f5e4d50 ->
../../sdc1

lrwxrwxrwx. 1 root root 10 Feb 28 11:02 cb1be19b-b546-4016-875a-2e626514e057 ->
../../sda2

lrwxrwxrwx. 1 root root 10 Mar 19 06:45 f4e88072-af9f-4a3b-aaa3-ee2fef96c56d ->
../../sdb1
```

```
cat /etc/fstab
#
# /etc/fstab
# Created by anaconda on Wed Feb 12 10:02:02 2025
#
# 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.
#
# After editing this file, run 'systemctl daemon-reload' to update systemd
# units generated from this file.
#
/dev/mapper/rhel-root / xfs defaults 0 0
UUID=cb1be19b-b546-4016-875a-2e626514e057 /boot xfs defaults 0 0
UUID=1970-1C3F /boot/efi vfat umask=0077,shortname=winnt 0 2
/dev/mapper/rhel-swap none swap defaults 0 0
/dev/disk/by-id/f4e88072-af9f-4a3b-aaa3-ee2fef96c56d /disk1 xfs defaults 0 1
/dev/disk/by-id/c5fc2403-d7a2-40db-8046-76ef3f5e4d50 /disk2 xfs defaults 0 1
```

b. Note if additional disks are mounted using the device name, such as:

```
/dev/sdb
/dev/sdc
```

c. If the devices remain mounted using the device names, **the mounts will need to be fixed after migration**, as the new devices will be **/dev/vdb**, **/dev/vdc**, and such. Additional details are listed in the [“Post migration”](#) section.

For example:

```
cat /etc/fstab
#
# /etc/fstab
# Created by anaconda on Wed Feb 12 10:02:02 2025
#
# 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.
#
# After editing this file, run 'systemctl daemon-reload' to update systemd
# units generated from this file.
```

```
#
/dev/mapper/rhel-root / xfs defaults 0 0
UUID=cb1be19b-b546-4016-875a-2e626514e057 /boot xfs defaults 0 0
UUID=1970-1C3F /boot/efi vfat umask=0077,shortname=winnt 0 2
/dev/mapper/rhel-swap none swap defaults 0 0
/dev/sdb1 /disk1 xfs defaults 0 1
/dev/sdc1 /disk2 xfs defaults 0 1
```

2. Note/save the networking connection device names and settings:

```
ip a
```

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
```

```
link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
```

```
inet 127.0.0.1/8 scope host lo
```

```
valid_lft forever preferred_lft forever
```

```
inet6 ::1/128 scope host
```

```
valid_lft forever preferred_lft forever
```

```
2: ens192: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
```

```
link/ether 00:50:56:9d:94:1e brd ff:ff:ff:ff:ff:ff
```

```
inet 10.3.23.225/24 brd 10.3.23.255 scope global dynamic noprefixroute ens192
```

```
valid_lft 5964sec preferred_lft 5964sec
```

```
inet6 fe80::250:56ff:fe9d:941e/64 scope link noprefixroute
```

```
valid_lft forever preferred_lft forever
```

```
nmcli device show
```

```
GENERAL.DEVICE: ens192
GENERAL.TYPE: ethernet
GENERAL.HWADDR: 00:50:56:9D:94:1E
GENERAL.MTU: 1500
GENERAL.STATE: 100 (connected)
GENERAL.CONNECTION: ens192
GENERAL.CON-PATH: /org/freedesktop/NetworkManager/ActiveConnection/1
WIRED-PROPERTIES.CARRIER: on
IP4.ADDRESS[1]: 10.3.23.225/24
IP4.GATEWAY: 10.3.23.1
IP4.ROUTE[1]: dst = 10.3.23.0/24, nh = 0.0.0.0, mt = 100
IP4.ROUTE[2]: dst = 0.0.0.0/0, nh = 10.3.23.1, mt = 100
IP4.DNS[1]: 10.3.23.2
```

```

IP4.DNS[2]:                10.3.23.3
IP4.DOMAIN[1]:             v0023.sau01.presales.hpecorp.net
IP6.ADDRESS[1]:           fe80::250:56ff:fe9d:941e/64
IP6.GATEWAY:              --
IP6.ROUTE[1]:             dst = fe80::/64, nh = ::, mt = 1024

```

```

GENERAL.DEVICE:           lo
GENERAL.TYPE:             loopback
GENERAL.HWADDR:           00:00:00:00:00:00
GENERAL.MTU:              65536
GENERAL.STATE:            10 (unmanaged)
GENERAL.CONNECTION:      --
GENERAL.CON-PATH:        --
IP4.ADDRESS[1]:           127.0.0.1/8
IP4.GATEWAY:              --
IP6.ADDRESS[1]:           ::1/128
IP6.GATEWAY:              --
IP6.ROUTE[1]:            dst = ::1/128, nh = ::, mt = 256

```

- a. After the migration, the networking device names may change and will need their names updated in **/etc/sysconfig/network-scripts/**, or previous settings need to be applied to the new devices using **nmcli**.
3. By default, the temporary root filesystem (initramfs) does not have the VirtIO drivers enabled in the OS. Follow these steps to ensure they are enabled:

- a. Verify that the drivers are not present; the following command should **not** return output:

```
lsinitrd /boot/initramfs-$(uname -r).img | grep -e virtio_scsi -e virtio_blk -e virtio_net
```

- l. If these drivers are already installed, then this step is complete, and you can skip to the next step.

- b. If the drivers are not present, to enable the required VirtIO drivers, add them to the configuration for initramfs:

```
echo 'add_drivers+=" virtio_scsi virtio_blk virtio_net "' | tee -a /etc/dracut.conf.d/virtio.conf
```

- c. After the configuration file is prepped, rebuild the initramfs:

```
dracut -f
```

- d. Once the rebuild is complete, check if the necessary VirtIO drivers are available using the following command:

```
lsinitrd /boot/initramfs-$(uname -r).img | grep -e virtio_scsi -e virtio_blk -e virtio_net
```

```

-rw-r--r--  1 root    root          9104 Jan 31  2022 usr/lib/modules/4.18.0-372.9.1.el8.x86_64/kernel/drivers/block/virtio_blk.ko.xz
-rw-r--r--  1 root    root        26172 Jan 31  2022 usr/lib/modules/4.18.0-372.9.1.el8.x86_64/kernel/drivers/net/virtio_net.ko.xz

```

```
-rw-r--r-- 1 root root 8644 Jan 31 2022 usr/lib/modules/4.18.0-372.9.1.el8.x86_64/kernel/drivers/scsi/virtio_scsi.ko.xz
```

4. If not already installed, install the qemu-guest-agent package:

```
dnf install qemu-guest-agent -y
```

- a. This may require an active Red Hat subscription and available repositories.
5. If one does not already exist, it is recommended to create a local user, which can be used to access the new virtual machine. In cases where the VM uses authentication from a remote source such as Active Directory, it may be unavailable if the network is not available immediately.
6. Eject any CD-ROM/DVD-ROMs that may be attached through VMware by selecting **Client Device** from the interface.




Figure 2. Eject CD

7. Shut down the source VM to ensure data consistency.

Migrating the VM on Red Hat

With the VM prepared, refer back to the section earlier in this document on creating Migration Plans. The HPE Morpheus VM Essentials Bulk Migration tool can easily and automatically migration prepared VMs to the HVM hypervisor.

Post migration Steps

After migration, there may be additional items that need to be resolved. The following are a few common items seen (some actions would require a privileged user account):

1. Network adapter name updates:
 - a. As mentioned in the [“Prepare the OS”](#) section, the networking device names may change and will need their names updated in `/etc/sysconfig/network-scripts/`, or previous settings need to be applied to the new devices using **nmcli**.
 - b. Example commands to rename the previous network manager connection and associate it with the new device, retaining the settings:

```
nmcli connection modify ens192 con-name enp3s0
nmcli connection modify enp3s0 connection.interface-name enp3s0
```

2. Device mounting when using device names that are not UUID or other static values:

- a. This information is used if disks are mounted using the device name, such as:

```
/dev/sdb
/dev/sdc
```

- b. If the devices remain mounted using the device names, **the mounts will need to be fixed after migration**, as the new devices will be `/dev/vdb`, `/dev/vdc`, and such. For example:

```
cat /etc/fstab
#
```

```
# /etc/fstab
# Created by anaconda on Wed Feb 12 10:02:02 2025
#
# 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.
#
# After editing this file, run 'systemctl daemon-reload' to update systemd
# units generated from this file.
#
/dev/mapper/rhel-root  /                               xfs      defaults      0 0
UUID=cb1be19b-b546-4016-875a-2e626514e057 /boot xfs      defaults      0 0
UUID=1970-1C3F        /boot/efi                          vfat
umask=0077,shortname=winnt 0 2
/dev/mapper/rhel-swap  none                               swap     defaults      0 0
/dev/sdb1              /disk1                             xfs      defaults      0 1
/dev/sdc1              /disk2                             xfs      defaults      0 1

```

- c. The new device names should align closely to the same, but first verify the available disks:

```
lsblk
```

```

NAME                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
vda                  8:0    0   16G  0 disk
├─vda1               8:1    0   600M  0 part /boot/efi
├─vda2               8:2    0    1G   0 part /boot
└─vda3               8:3    0  14.4G  0 part
   ├─rhel-root       253:0   0  12.8G  0 lvm  /
   └─rhel-swap       253:1   0   1.6G  0 lvm  [SWAP]
vdb                  8:16   0  100G  0 disk
└─vdb1               8:17   0  100G  0 part /disk1
vdc                  8:32   0  101G  0 disk
└─vdc1               8:33   0  101G  0 part /disk2

```

- d. Prior to migration the device names were **sdb** and **sdc** but are now **vdb** and **vdc**, respectively. However, they are not the same; be sure to map them accordingly.

Additionally, output saved from running **ls -lah /dev/disk/by-uuid** while executing the steps list in Prepare the source VM section can help identify the disks, as the UUID should not change. For example:

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

```
Premigration:
```

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

```

total 0
drwxr-xr-x. 2 root root 180 Mar 19 06:45 .
drwxr-xr-x. 8 root root 160 Feb 28 11:02 ..
lrwxrwxrwx. 1 root root  10 Feb 28 11:02 1970-1C3F -> ../../sda1
lrwxrwxrwx. 1 root root   9 Feb 28 11:02 2022-06-28-07-51-01-00 -> ../../sr0

```

```

lrwxrwxrwx. 1 root root 10 Feb 28 11:02 44e3e59c-fc46-4c1a-808c-e2d31b8fc57b ->
.././dm-1
lrwxrwxrwx. 1 root root 10 Feb 28 11:02 768394c6-7eca-48ad-a711-c93588e1b4f6 ->
.././dm-0
lrwxrwxrwx. 1 root root 10 Mar 19 06:45 c5fc2403-d7a2-40db-8046-76ef3f5e4d50 ->
.././sdc1
lrwxrwxrwx. 1 root root 10 Feb 28 11:02 cb1be19b-b546-4016-875a-2e626514e057 ->
.././sda2
lrwxrwxrwx. 1 root root 10 Mar 19 06:45 f4e88072-af9f-4a3b-aaa3-ee2fef96c56d ->
.././sdb1

```

Post migration:

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

```

total 0
drwxr-xr-x. 2 root root 180 Feb 28 11:45 .
drwxr-xr-x. 8 root root 160 Feb 28 11:45 ..
lrwxrwxrwx. 1 root root 10 Feb 28 11:45 1970-1C3F -> .././vda1
lrwxrwxrwx. 1 root root 10 Feb 28 11:45 44e3e59c-fc46-4c1a-808c-e2d31b8fc57b ->
.././dm-1
lrwxrwxrwx. 1 root root 10 Feb 28 11:45 768394c6-7eca-48ad-a711-c93588e1b4f6 ->
.././dm-0
lrwxrwxrwx. 1 root root 10 Feb 28 11:45 c5fc2403-d7a2-40db-8046-76ef3f5e4d50 ->
.././vdc1
lrwxrwxrwx. 1 root root 10 Feb 28 11:45 cb1be19b-b546-4016-875a-2e626514e057 ->
.././vda2
lrwxrwxrwx. 1 root root 10 Feb 28 11:45 f4e88072-af9f-4a3b-aaa3-ee2fef96c56d ->
.././vdb1

```

- e. Once all device names are sorted, update the **/etc/fstab** with the appropriate device names and partition, for example:

```

# /etc/fstab
# Created by anaconda on Wed Feb 12 10:02:02 2025
#
# 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.
#
# After editing this file, run 'systemctl daemon-reload' to update systemd
# units generated from this file.
#
/dev/mapper/rhel-root / xfs defaults 0 0
UUID=cb1be19b-b546-4016-875a-2e626514e057 /boot xfs defaults 0 0
UUID=1970-1C3F /boot/efi vfat umask=0077,shortname=winnt 0 2
/dev/mapper/rhel-swap none swap defaults 0 0

```

```

/dev/vdb1          /disk1          xfs    defaults    0 1
/dev/vdc1          /disk2          xfs    defaults    0 1

```

f. Once the mounts are updated correctly, use the following to mount them:

```

mount -a
systemctl daemon-reload

```

Using `df -h` or other tools should now show the partitions mounted to the locations specified:

```

df -h

```

Filesystem	Size	Used	Avail	Use%	Mounted on
devtmpfs	888M	0	888M	0%	/dev
tmpfs	906M	0	906M	0%	/dev/shm
tmpfs	906M	8.7M	898M	1%	/run
tmpfs	906M	0	906M	0%	/sys/fs/cgroup
/dev/mapper/rhel-root	13G	2.2G	11G	17%	/
/dev/vda2	1014M	236M	779M	24%	/boot
/dev/vda1	599M	5.9M	594M	1%	/boot/efi
tmpfs	182M	0	182M	0%	/run/user/0
/dev/vdb1	100G	1.1G	99G	2%	/disk1
/dev/vdc1	101G	754M	101G	1%	/disk2

Rollback plan

If the VM fails to come up on HPE Morpheus VM Essentials Software, shut down the VM on the software hypervisor. Power on the source VM either from vCenter or the software user interface.

Windows Server 2022 (or later)

Supported versions and assumptions

- Windows Server 2022 or later
- The VM is configured to boot as EFI or BIOS
- The VM has Virtualization Based Security (VBS) enabled
- The VM has Secure boot enabled
- The VM has the Trusted Platform Module (TPM) disabled
- The VM has access to the recovery partition or a Windows Recovery Environment (RE) disc or installation media is available
- Some of the above assumptions may not be a requirement but were used in the testing process

Prepare the OS

Inject VirtIO drivers:

1. Restart the VMware Windows guest in the [Windows Recovery Environment \(Windows RE\)](#):
There are many ways, but this was used in this example:

Settings > Update & Security > Recovery > Restart Now button under Advanced startup.

2. Once in the Windows RE, choose your keyboard layout if prompted.
3. Click the **Troubleshoot** option.

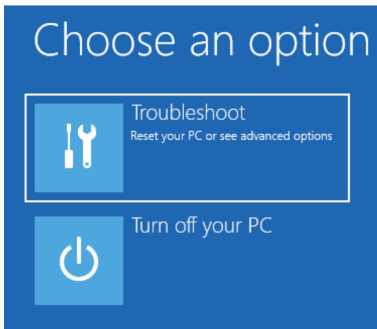


Figure 3. Troubleshoot option

4. Click the **Command Prompt** option, which will launch a command prompt.

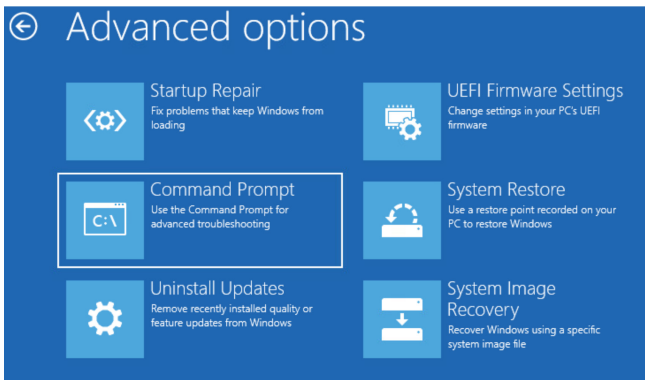


Figure 4. Choose Command Prompt

5. Enter diskpart using this command:

diskpart

6. Use the following command to see the disks that are available.

list disk

7. **(Optional)** If you receive the message **There are no fixed disks to show** or the disk with the Windows OS is not listed when using the previous command, follow these steps to mount the drives; otherwise, skip to the next step:

DISKPART> list disk

There are no fixed disks to show

- a. In vCenter, mount the VMware Tools™ installer to the VM using the **Install VMware Tools...** option:

- I. In this example, we used:

Actions > Guest OS > Install VMware Tools...

- b. While still in diskpart, use the following command to find the CD-ROM/DVD-ROM drive letter:

list volume

- c. Note the CD-ROM drive letter; in this example, we'll use the letter **D**.

DISKPART> list volume

Volume ###	Ltr	Label	Fs	Type
-----	---	-----	----	-----

Volume 0 D VMware Tool CDFS DVD-ROM

d. Exit diskpart by entering:

exit

e. Enter the following command to load the storage drivers to be able to see the hard disks:

drvload "D:\Program Files\VMware\VMware Tools\Drivers\pvscsi\Win10\amd64\pvscsi.inf"

l. Note that the previous mentions win10 as the operating system for the driver, but choose the most appropriate or latest available.

f. Once the driver has been loaded, use the following commands to verify the disks are mounted:

diskpart

list disk

X:\Windows\system32>diskpart

Microsoft DiskPart version 10.0.17763.1

Copyright (C) Microsoft Corporation.

On computer: WIN-ROCNE078V4K

DISKPART> list disk

Disk ###	Status	Size	Free	Dyn	Gpt
-----	-----	-----	-----	---	---
Disk 0	Online	40 GB	1024 KB		*
Disk 1	Online	1024 MB	1024 MB		

g. With the disks mounted, the partitions/volumes will need drive letters, while still in diskpart, enter the following command to list the volumes:

list volume

DISKPART> list volume

Volume ###	Ltr	Label	Fs	Type	Size	Status	Info
-----	---	-----	-----	-----	-----	-----	-----
Volume 0	D	VMware Tool	CDFS	DVD-ROM	124 MB	Healthy	
Volume 1			NTFS	Partition	39 GB	Healthy	
Volume 2			FAT32	Partition	100 MB	Healthy	Hidden
Volume 3			NTFS	Partition	524 MB	Healthy	Hidden

h. As seen previously, our primary Windows volume is **Volume 1** (the big drive), which we'll use in the following examples.

i. Set a drive letter for **volume 1** using the following commands:

select volume 1
assign letter=c

```
DISKPART> select volume 1
```

Volume 1 is the selected volume.

```
DISKPART> assign letter=c
```

DiskPart successfully assigned the drive letter or mount point.

j. The volume letter can be verified using:

list volume

```
DISKPART> list volume
```

Volume ###	Ltr	Label	Fs	Type	Size	Status	Info
Volume 0	D	VMware Tool	CDFS	DVD-ROM	124 MB	Healthy	
* Volume 1	C		NTFS	Partition	39 GB	Healthy	
Volume 2			FAT32	Partition	100 MB	Healthy	Hidden
Volume 3			NTFS	Partition	524 MB	Healthy	Hidden

k. Exit diskpart:

exit

l. In vCenter, unmount the VMware Tools installer to the VM using the Unmount VMware Tools Installer option. In this example, we used **Actions > Guest OS > Unmount VMware Tools Installer**.

8. Download either the **latest** or **stable** VirtIO drivers **ISO** on your local workstation:

Download link: github.com/virtio-win/virtio-win-pkg-scripts/blob/master/README.md

a. The **latest** is recommended, but most often they are the same.

b. Latest: [Latest VirtIO ISO](#)

c. Stable: [Stable VirtIO ISO](#)

9. Add the VirtIO drivers ISO to vCenter, either as an upload to a datastore or to the Content Library.

10. Mount/attach the ISO using a CD-ROM device on the virtual machine and make sure it is **Connected**.




Figure 5. Attach VirtIO ISO

a. This will replace the mounted VMware Tools installer, if it was needed in the Optional step previously.

11. Enter the following command to inject the VirtIO storage drivers into the Windows [boot-start drivers](#) from the mounted ISO:

```
dism /image:C:\ /add-driver:D:\viosstor\2k22\amd64\viosstor.inf
```

```
dism /image:C:\ /add-driver:D:\vioscsi\2k22\amd64\vioscsi.inf
```

- a. If the OS volume letter assigned was not **C**, then use the appropriate letter.
- b. If the CD-ROM/DVD-ROM volume letter assigned was not **D**, then use the appropriate letter.
- c. In the previous example, **2k22** is used, but select the appropriate OS version as needed.
- d. If successful:

```
X:\Windows\system32>dism /image:C:\ /add-driver:D:\viosstor\2k22\amd64\viosstor.inf
```

```
Deployment Image Servicing and Management tool
```

```
Version: 10.0.17763.1
```

```
Image Version: 10.0.17763.771
```

```
Found 1 driver package(s) to install.
```

```
Installing 1 of 1 - oem21.inf: The driver package was successfully installed.
```

```
The operation completed successfully.
```

12. The command prompt can be closed by clicking the red X:



Figure 6. Close Command Prompt Window

13. Click **Continue** to boot up the operating system again.

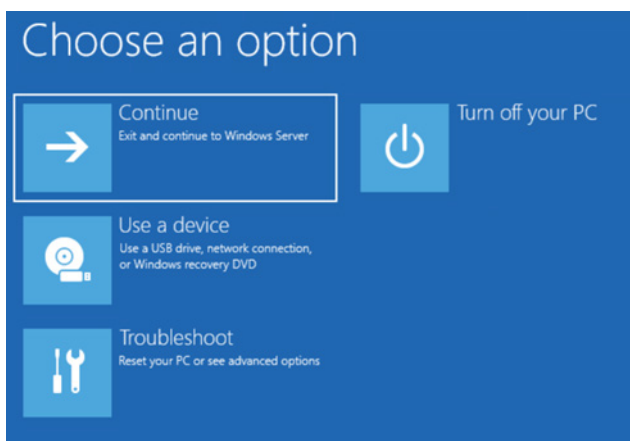


Figure 7. Continue OS Boot process

14. Once logged back into the OS, install the VirtIO drivers in the OS using the ISO, still mounted, by navigating to the CD-ROM/DVD-ROM drive:

- a. Install the **virtio-win-gt-x64.msi** and **virtio-win-guest-tools.exe** from the root of the ISO, keeping the default selections.

b. Example commands to do the same as above, assuming the CD-ROM/DVD-ROM driver letter is **D:**

D:\virtio-win-gt-x64.msi /q

D:\virtio-win-guest-tools.exe /q

15. Unmount any ISOs if it is still attached to the VM, either by ejecting through Windows or setting the CD-ROM/DVD-ROM drive to **Client Device** in vCenter.

16. Select an option, either:

a. Turn off the VM (recommended).

I. When turning off the VM, it can be turned back on automatically if the following setting is configured on the VMware Cloud™:

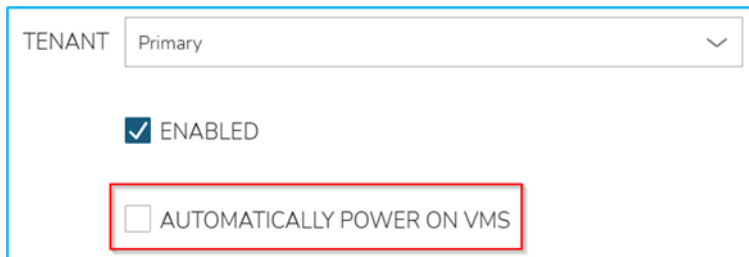


Figure 8. Uncheck Automatically power on VMs

II. Either:

- Be sure to shut down the VM through the HPE Morpheus VM Essentials Software UI (so it won't turn on the VM automatically).
- Turn off the setting previously mentioned if shutting down from the guest or VMware native console.

b. Keep the VM online:

- I. The VM will be quiesced before exporting, but note that a **diff of data will not be performed** once the migration is complete.
- II. This can cause an issue if the same IP is used on the network, two systems may be online with the same IP.

Migrating the VM on Windows

With the VM prepared, refer back to the section earlier in this document on creating Migration Plans. The HPE Morpheus VM Essentials Bulk Migration tool can easily and automatically have the migration prepare VMs to move to the HVM hypervisor.

Rollback plan

If the VM fails to come up on HPE Morpheus VM Essentials Software, shut down the VM on the software hypervisor. Power on the source VM either from vCenter or the software user interface.

Troubleshooting

HPE Morpheus VM Essentials Software provides an option to export logs from the **Health** menu option. To export the logs, navigate to **Administration > Health > Morpheus Logs > Options > export**. These logs can help the user troubleshoot and investigate issues faced while performing the VM image conversion tasks.

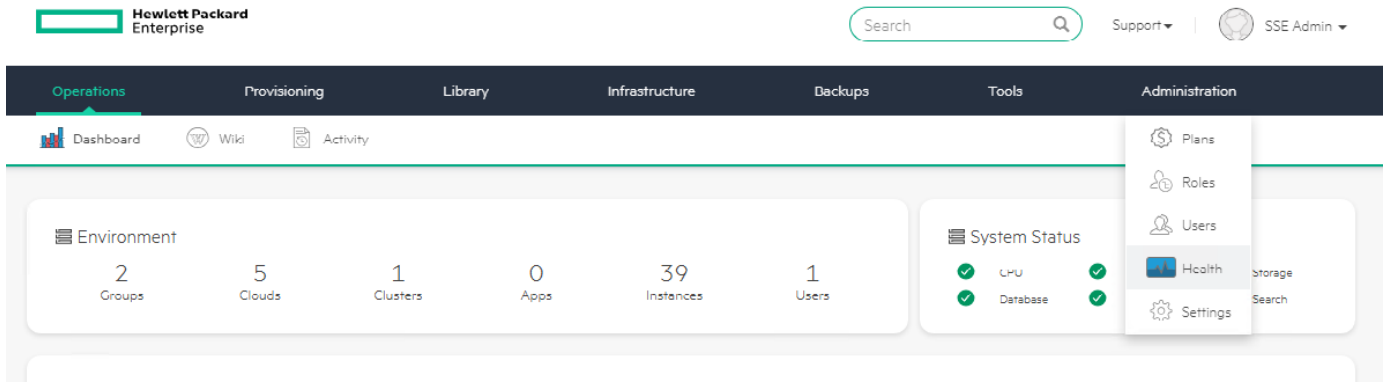


Figure 9. Health Menu Option

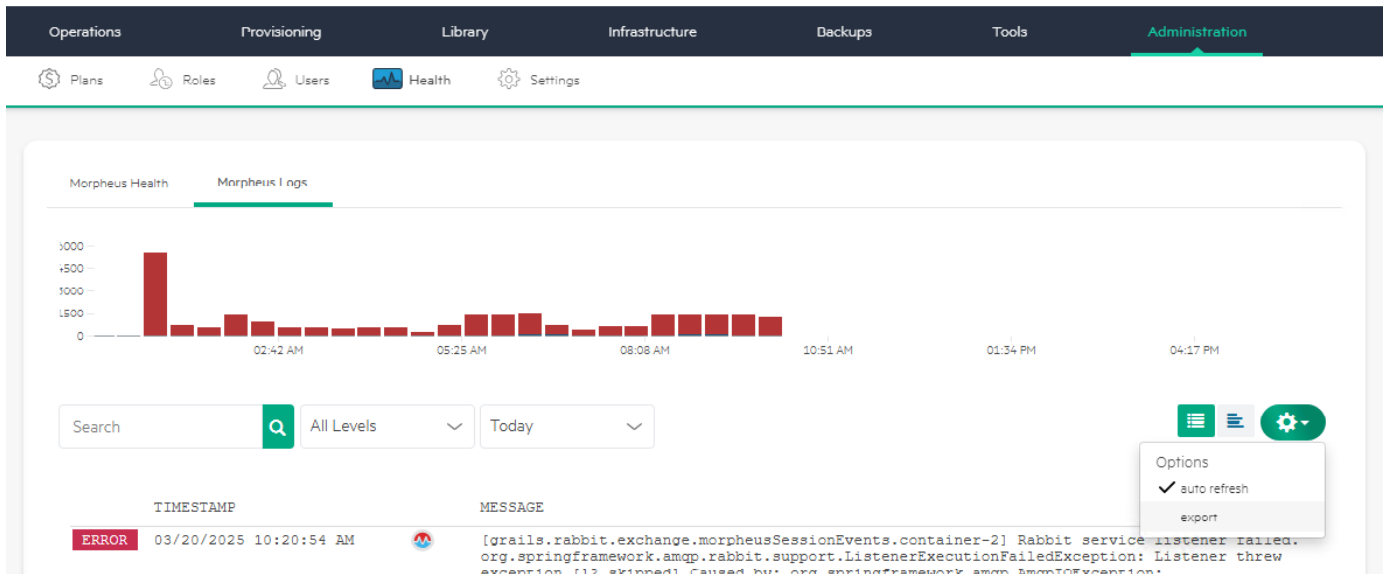


Figure 10. View Morpheus Logs

Conclusion

In this document, we have outlined the steps to migrate VMs from VMware-based environments to HPE Morpheus VM Essentials Software. See the [Reference links](#) section for additional information about configuring and optimizing the migrated VMs and to take advantage of the full set of capabilities provided by HPE Morpheus VM Essentials Software.

Appendix

Procedure to create an external datastore for images and VMs

HPE Morpheus VM Essentials Software cluster can be configured with HPE Alletra Storage MP B10000, either using iSCSI or FC protocol. Shared storage can be used for storing VM images and VM datastores.

All the software hosts in the cluster must be configured with HPE Alletra Storage MP B10000 for shared storage access.

Prerequisites

- Configuration best practices are followed to avoid any single point of failure.
- Create storage volumes on HPE Alletra Storage MP B10000 array.
- Volumes are exported to the software hosts of the cluster.
- Multipathing is configured.

Procedure

From the software UI, navigate to **Infrastructure > Clusters > HPE VME Cluster > Storage > Add**, and follow the instructions to **Add the Datastore**.

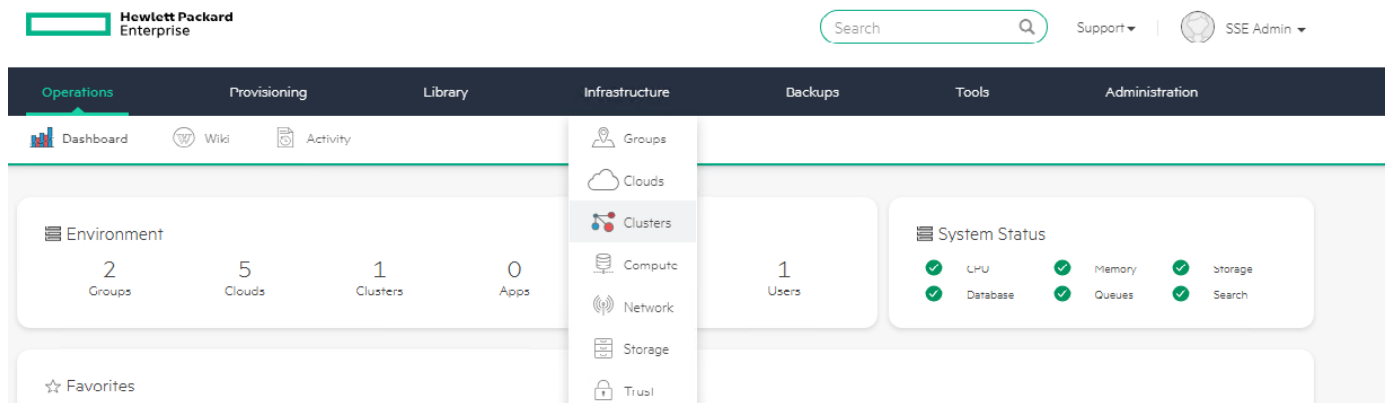


Figure 11. Clusters view selection screen

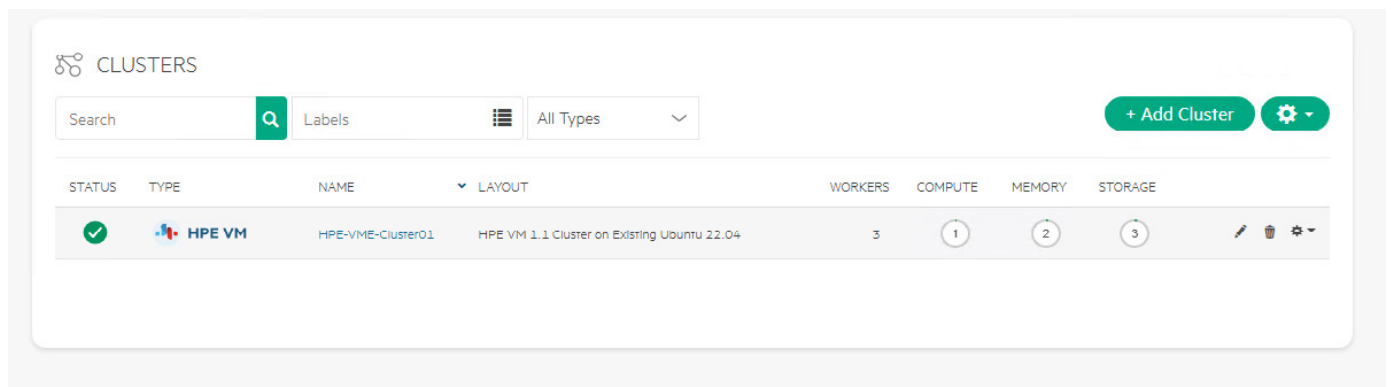


Figure 12. Select the cluster

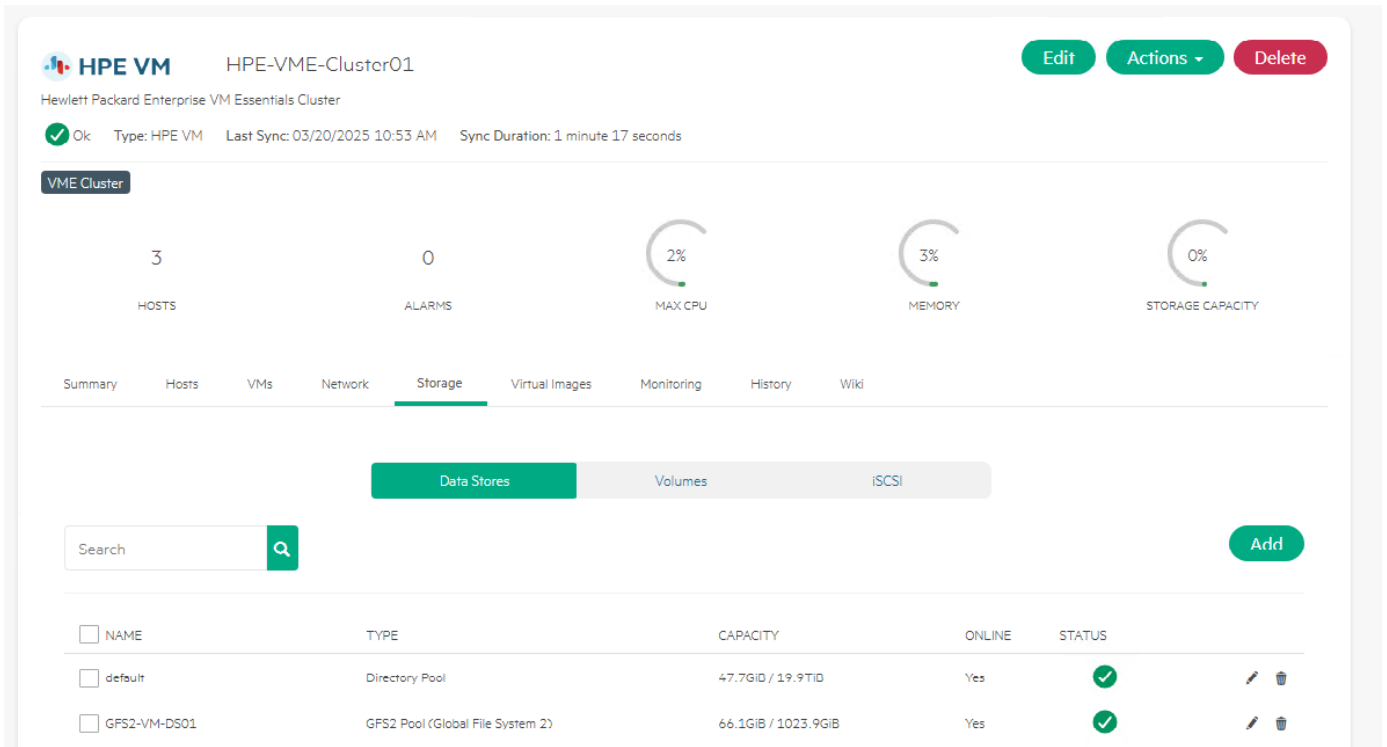


Figure 13. Storage details screen

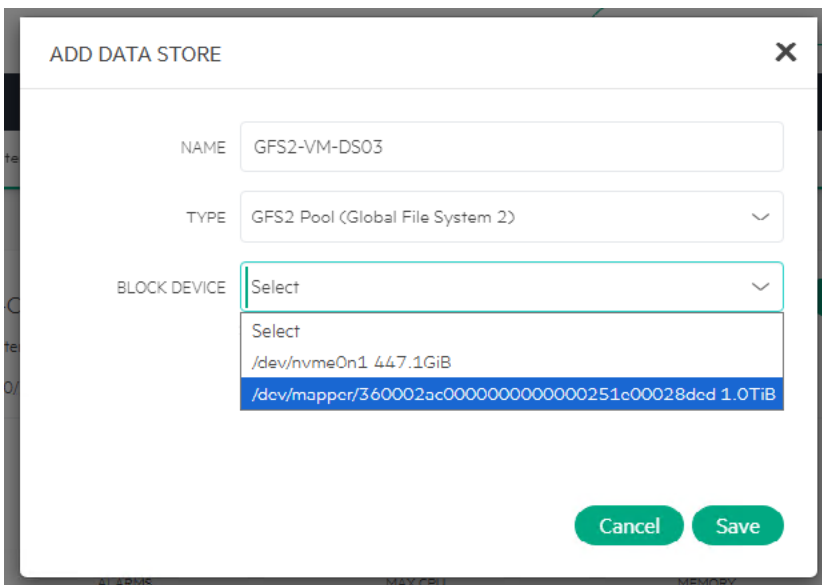


Figure 14. Add data store

Reference links

- [HPE Morpheus VM Essentials Software](#)
- [HPE Morpheus VM Essentials Software demo](#)
- [HPE Morpheus VM Essentials Software QuickSpecs](#)
- [HPE Morpheus VM Essentials Software Documentation](#)

Learn more at

[HPE.com/vmessentials](https://hpe.com/vmessentials)

Visit [HPE.com](https://hpe.com)

[Chat now](#)

© Copyright 2025 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Active Directory, Microsoft, Windows, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. VMware Cloud, VMware ESXi, VMware Tools, VMware vCenter, and VMware are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All third-party marks are property of their respective owners.

a50013873ENW

HEWLETT PACKARD ENTERPRISE

hpe.com

