



# **Intel<sup>®</sup> Volume Management Device (Intel<sup>®</sup> VMD) Driver for VMware\* ESXi Red Hat\* Enterprise Linux\* Virtual Machines Direct Assign**

**User Guide**

---

***Revision 003***

***March 2024***



You may not use or facilitate the use of this document in connection with any infringement or other legal analysis. You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

All product plans and roadmaps are subject to change without notice.

The products described may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at [intel.com](https://www.intel.com).

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Altering clock frequency, voltage, or memory interface speeds may void any product warranties and reduce stability, security, performance, and life of the processor and other components. Intel has not validated processor running memory above Plan-Of-Record (POR) speed. DRAM/DIMM devices should support desired speed, check with DRAM/DIMM vendors for details. System manufacturers are responsible for all validation and assume the risk of any stability, security, performance, or other functional issues resulting from such alterations.

© Intel Corporation. Intel, the Intel logo, Xeon®, and other Intel marks are trademarks of Intel Corporation or its subsidiaries.

\*Other names and brands may be claimed as the property of others.

Copyright© 2023-2024 Intel Corporation. All rights reserved.

# Contents

---

1	Overview .....	5
	1.1 Test Specifications .....	5
	1.2 Test Scenarios.....	5
2	Enable VMD Direct Assign in the BIOS Setup Menu .....	6
3	Find VMD Bus: Device.Function (B:D.F) Designator for Drives .....	7
4	Install ESXi*, Create and Configure the VM.....	9
5	Adding VMD PCIe Devices to the VM .....	12
6	Linux* OS Installation with RHEL* Inbox VMD Driver .....	14
7	Secure Boot using Linux* Inbox Certificate .....	15
8	VMD Replacement Driver Details.....	16
9	Add Intel® Replacement Driver to Existing VM.....	17
10	Secure Boot using Intel® Replacement Certificate .....	19
11	Troubleshooting.....	21
	11.1 Direct Assigned NVMe's Do Not Display.....	21
	11.2 VM Does Not Boot with Secure Boot Enabled .....	21
	11.3 Gen1/2 Xeon System Behavior and Recommendations .....	21
	11.4 On VM Startup, Error Shows with Failure to Power On due to Wrong Memory Size .....	21
	11.5 When Using a Replacement Driver in the VM, NVMe* Devices are Not Visible	22
	11.6 Error When Installing RHEL OS to VM .....	22



## Revision History

Revision Number	Description	Revision Date
001	• Initial Release	September 2023
002	• Updated Chapter <a href="#">5</a>	January 2024
003	• Updated Chapter <a href="#">11</a>	March 2024

§§

# 1 Overview

---

This guide describes setup and configuration of system BIOS, ESXi\* Hypervisor, and Red Hat\* Enterprise Linux\* (RHEL) Virtual Machine (VM) drive assignment using Direct Assign (DA).

## 1.1 Test Specifications

- ESXi\* 8.0U1 GA Release
- ESXi\* VMD Inbox Driver 3.0.0.1010
- RHEL\* 9.1 GA Release with inbox VMD Driver and certificate
- Intel® replacement driver and certificate for RHEL\* 9.1
- Intel® Xeon® Scalable Platform Generation 4 ESXi\* Host

## 1.2 Test Scenarios

- DA using ESXi\* and RHEL\* inbox VMD drivers.
- DA with Secure Boot using ESXi\* and RHEL\* VMD inbox drivers and certificate.
- DA Using Intel replacement VMD driver.
- DA with Secure Boot using Intel replacement VMD driver and certificate.

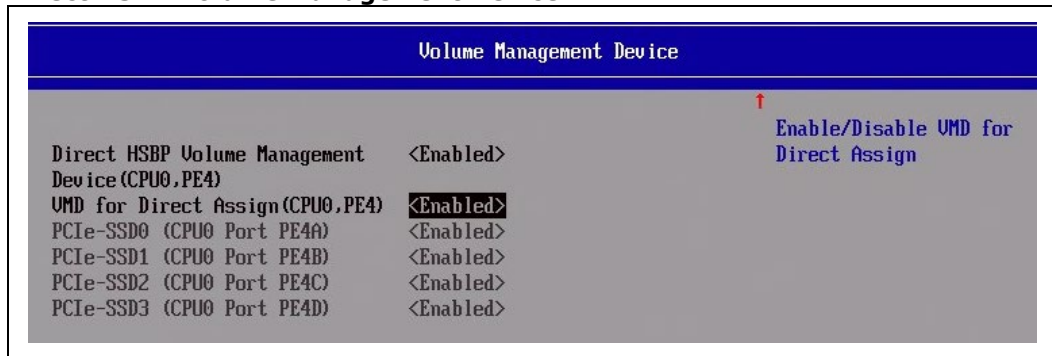
The following chapters assume a working knowledge of ESXi\* and RHEL\* installation and administration. Screenshots are provided to help illustrate ESXi\* and RHEL\* configuration settings.



## 2 Enable VMD Direct Assign in the BIOS Setup Menu

On Intel® Xeon® Scalable Platform Generation 3 platforms and newer, enable **VMD for Direct Assign** within the BIOS setup menu for the VMD domain where NVMe(s) are attached. For the System Under Test (SUT), menu navigation follows:

**Advanced > Integrated IO Configuration > Volume Management Device > Direct HSBP Volume Management Device**



For the SUT, the above BIOS screenshot enables Direct Assign capability for VMD (CPU0, PE4), which relates to drive bay slots 1-4 in this system. BIOS menu entries may differ based on BIOS vendor or OEM system architecture. Some OEMs may mask certain BIOS settings, including VMD options. If uncertain as to where to find the VMD-specific menus, consult with the OEM or Intel® representative.

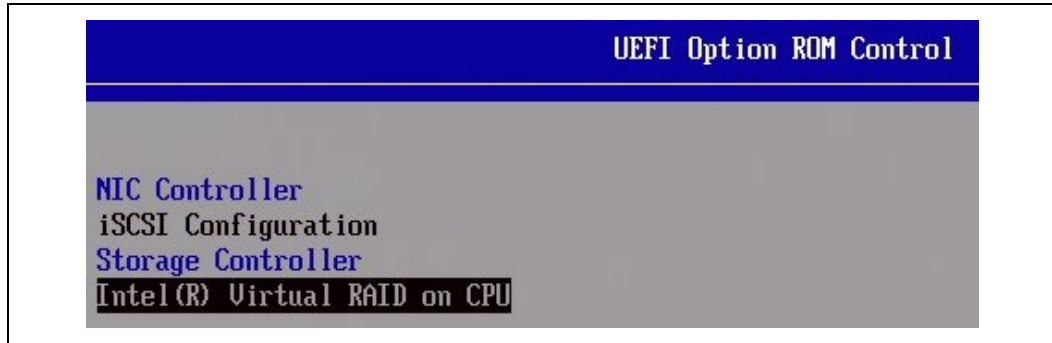
Once Direct Assign has been enabled for the target NVMe\* drive, reboot the system.




### 3 Find VMD Bus: Device.Function (B:D.F) Designator for Drives


To discover which VMD controller B:D.F owns the NVMe\* device(s) required for VM assignment, enter the BIOS Advanced menu as follows:

**Advanced > PCI Configuration > UEFI Option ROM Control > Intel Virtual RAID on CPU > All Intel VMD Controllers**





The above example indicates this system has two NVMe\* devices behind **VMD4**.



To assign both of these NVMe's to the VM, confirm both drives hold the same **VMD Bus:Device.Function**. In this case both show **62:00.5**.

PCI Bus:Device.Function is not used for this procedure.

**Note:** All NVMe\* drives owned by a specific VMD Controller (in this case VMD4) will be Direct Assigned as a group. It is not possible at this writing to split Controller/NVMe\* assignments.






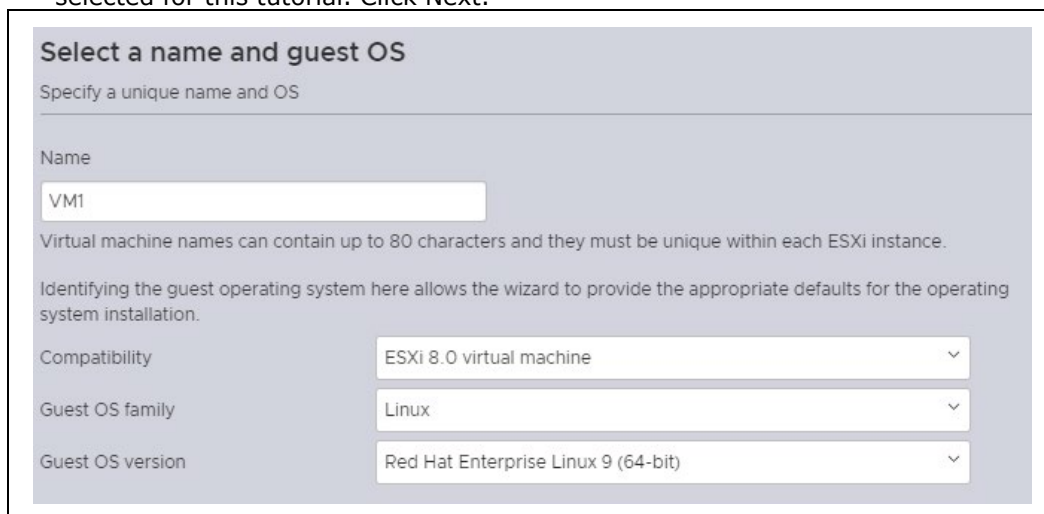
## 4 Install ESXi\*, Create and Configure the VM

1. Install VMware\* ESXi version 8.0U1 GA on the target ESXi\* Host. Do not install ESXi\* on a disk that will be direct assigned.
2. With the ESXi\* server up and running, open a web browser to the VMware\* ESXi Host Client at [https://\[ESXi IP\]/ui](https://[ESXi IP]/ui) and log in. ESXi\* IP is the network address of the server. Log in with credentials assigned at the time of ESXi\* installation.
3. Use an existing datastore or create a new datastore to assign to the VM. Ensure the datastore does not reside on a disk that will be direct assigned.
4. Select **Virtual Machines** from the Navigator pane. Select **Create / Register VM**.

**Note:** When creating the VM, it is automatically registered.

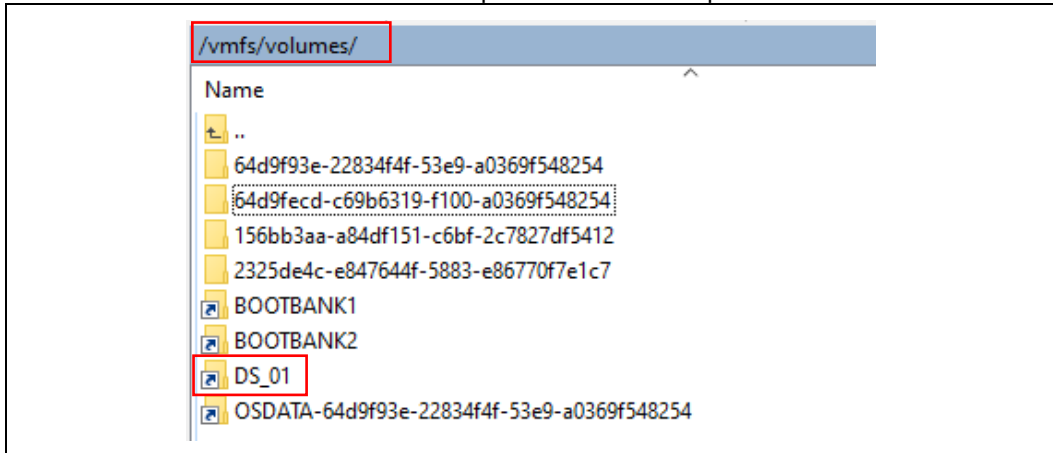


5. Create a new virtual machine. Click Next.
6. Name the VM. Keep the default Compatibility. Select *Linux* from the Guest OS family drop down menu. Select the version of RHEL\* to be installed from the Guest OS version drop down menu. *Red Hat\* Enterprise Linux\* 9 (64-bit)* is selected for this tutorial. Click Next.

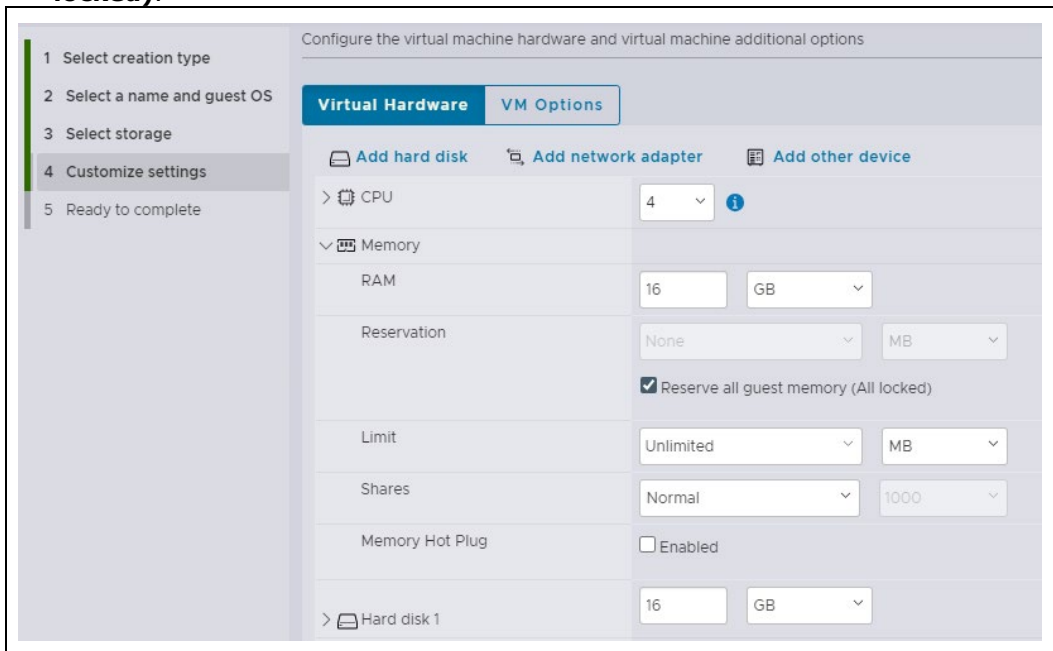


7. On the Select storage page, select the datastore created previously; DS\_01 is the assigned datastore name for this installation. Click Next.
8. To prepare for the VM OS installation, upload the RHEL\* 9.1 image (.ISO) to the datastore folder intended for VM OS installation.

- 9. Instructions for uploading the image:
  - a. Enable SSH and ESXi\* shell from the ESXi\* Troubleshooting Mode Options menu.
  - b. Use WinSCP\* to copy RHEL\* image to the datastore folder.
  - c. RHEL\* OS installation will take place at a later step.

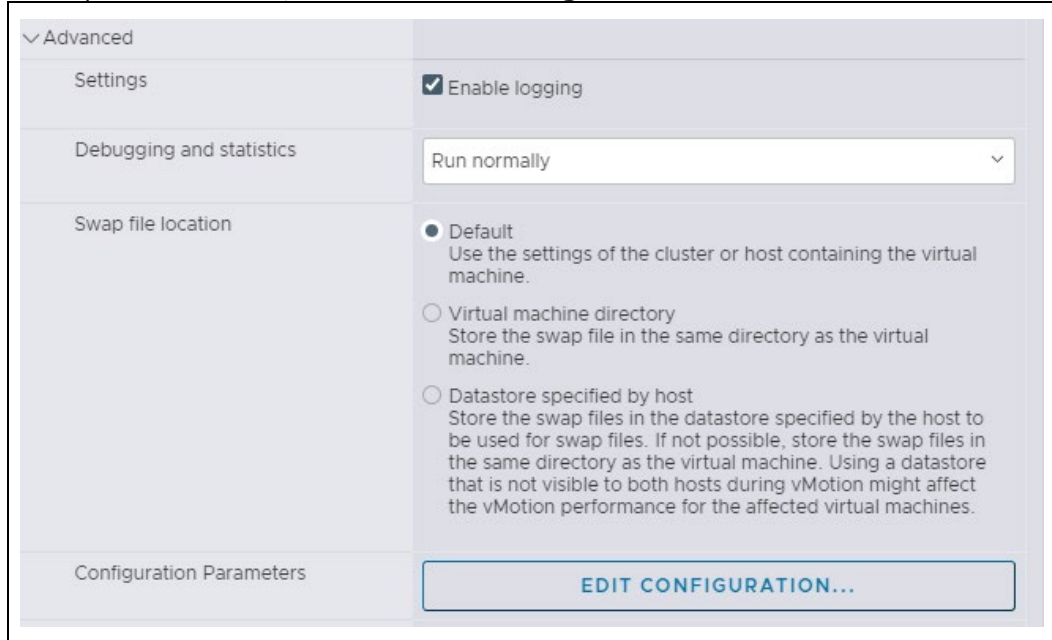


- 10. From New Virtual Machine menu, select **Customize settings > Virtual Hardware > CD/DVD Drive 1**.
- 11. Since we are installing from a .ISO image, select CD/DVD Drive 1 > Datastore ISO file from the dropdown menu. Highlight and select the .ISO previously uploaded to the datastore.
- 12. Set number of CPU's at least equal to the number of VMD domains which will be assigned to the VM. If only one virtual CPU is selected, but more than one domain is direct assigned, only one domain's drives will be visible.
- 13. Set Memory at 16 GB RAM and check the box **Reserve all guest memory (All locked)**.




**Install ESXi\*, Create and Configure the VM**

14. Select VM Options Button
15. Boot Options > unselect Whether or not to enable UEFI secure boot for this VM. We will address Secure Boot in later sections.
16. Expand **Advanced**, and select **Edit Configuration**



17. Add Parameter: *iavmd.direct\_assign* Value: *1*




18. Click **OK > Next > Finish**.



## 5 Adding VMD PCIe Devices to the VM

- The following steps modify ESXi\* to manage VM resource requirements.
- Add the following lines to ESXi\* host `/etc/vmware/passthru.map` file (using SSH, WinSCP\*, etc.) and save.
 

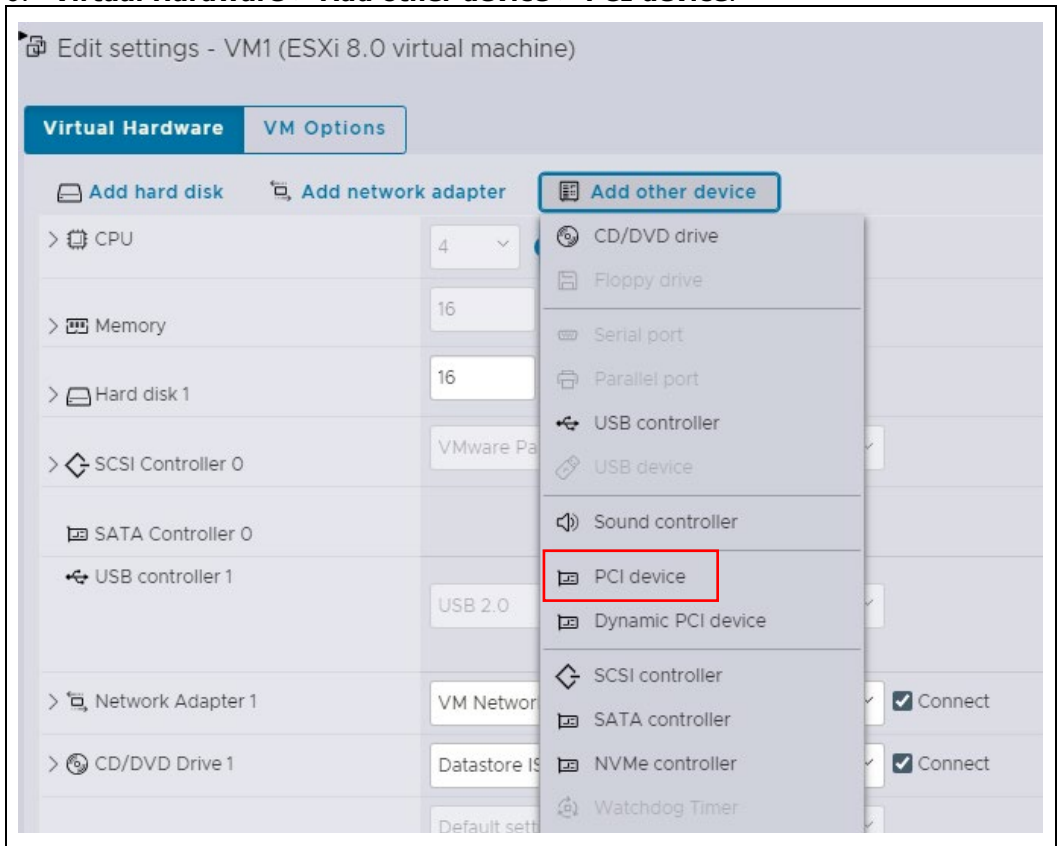
```
# Intel® Xeon® Scalable Platform Gen 3_4 VMD Controller Device ID
8086 28c0 d3d0 default
```
- In the ESXi\* Host Client Navigator pane, select Host > Manage. From the Hardware Menu > PCI Devices, enable VMD as passthrough in the Host by searching for the VMD controller B:D.F that was identified earlier for passthrough drives. Select the device and click Toggle passthrough from the top left.



- Note device status changed to *Active*.

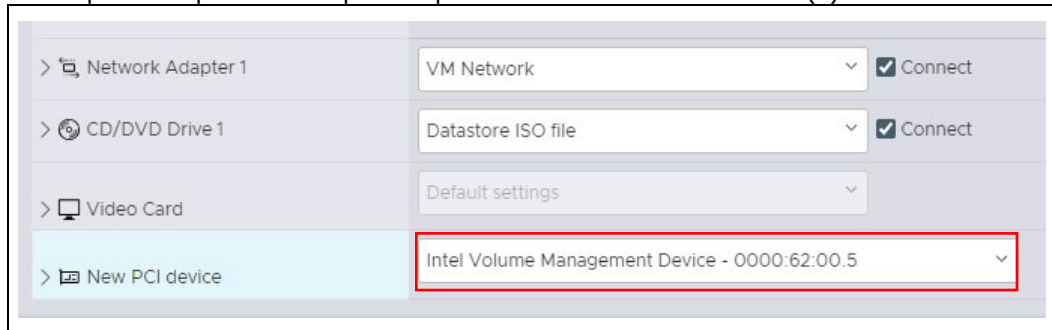


- Before powering on the VM, Select the VM and choose Actions > Edit settings.
- Virtual Hardware > Add other device > PCI device.**



**Adding VMD PCIe Devices to the VM**

7. After choosing to add the PCI device, a New PCI device entry will display at the bottom of the device list exposing the VMD controller as shown below. Click Save.
8. Repeat the previous steps if required for additional PCI device(s).



9. After the PCI device(s) is added, it may be necessary to reboot the ESXi host to ensure NVMe drive visibility. Behavior can vary depending on existing system configuration.



## 6 Linux\* OS Installation with RHEL\* Inbox VMD Driver

1. Select the VM from the Virtual Machines tab in Navigator. Select Power on.
2. After VM successfully powers on, from the Host Client UI, select Console > Open browser console in new window. At the RHEL\* installation options screen, Install Red Hat\* Enterprise Linux using preferred options.
3. After installation is complete, Reboot the VM.
4. Verify the Direct Assign/passthrough NVMe\* devices are present: **nvme list**, **lsblk**, or **ls -la /sys/block**. If NVMe\* devices are not present, refer to [Chapter 11](#).

```
[root@localhost ~]# lsblk
NAME                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINTS
sda                  8:0    0   16G  0 disk
├─sda1                8:1    0   600M  0 part /boot/efi
├─sda2                8:2    0    1G   0 part /boot
└─sda3                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]
sr0                  11:0    1  1024M  0 rom
nvme0n1              259:0    0    7T   0 disk
nvme1n1              259:1    0    7T   0 disk
[root@localhost ~]#
```

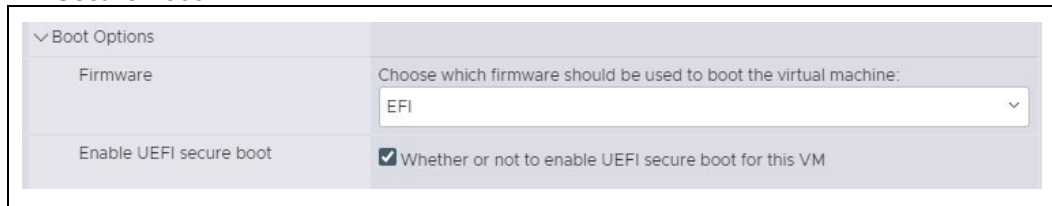
```
[root@localhost ~]# nvme list
Node          Generic          SN              Model
Namespace    Usage            Format           FW Rev
-----
/dev/nvme1n1  /dev/ng1n1      PHAB949200697P6GGN INTEL SSDPF2KX076T9
1             7.68 TB / 7.68 TB 512 B + 0 B 2CV10013
/dev/nvme0n1  /dev/ng0n1      PHAB9492001W7P6GGN INTEL SSDPF2KX076T9
1             7.68 TB / 7.68 TB 512 B + 0 B 2CV10013
```

§§

## 7 Secure Boot using Linux\* Inbox Certificate

---

1. Complete all previous steps without Secure Boot, confirm DA devices are visible, then perform the following to configure Secure Boot for the VM.
2. Power off the VM. Go to Edit Settings. Under VM Options > **Boot Options** > Enable UEFI secure boot. Check the box for **Whether or not to enable UEFI secure boot for this VM**. Click Save. Power On VM. VM is now configured for Secure Boot.



The screenshot shows the 'Boot Options' section of a virtual machine's settings. It is divided into two rows. The first row is for 'Firmware', with a dropdown menu set to 'EFI'. The second row is for 'Enable UEFI secure boot', with a checked checkbox labeled 'Whether or not to enable UEFI secure boot for this VM'.

▼ Boot Options	
Firmware	Choose which firmware should be used to boot the virtual machine: EFI
Enable UEFI secure boot	<input checked="" type="checkbox"/> Whether or not to enable UEFI secure boot for this VM



## 8 VMD Replacement Driver Details

---

In Chapters [1-7](#), Direct Assign and Secure Boot tests use RHEL\* GA inbox drivers and certificates. In [Chapter 9](#), Replacement Driver installation requiring additional configurations for DA is explained. Supplemental information is provided below.

**Intel® Xeon® Generation 3 and above** – RHEL\* Inbox drivers and certificates are signed and functional. No replacement driver or replacement certificate is required unless the user requires an optimization provided by the replacement driver.

### Replacement Driver package contents

The signed replacement driver package typically includes a replacement driver, replacement certificate, Driver Update Disk iso, and source files:

- **kmod-iavmd-x.x.x.xxxx-rhel\_9x.x86\_64.rpm** – Replacement Driver
- **Intel® VMD Linux Driver.der** – Replacement Certificate for Secure Boot
- **iavmd- x.x.x.xxxx-rhel\_9x-dud.iso** – Install Replacement Driver at Boot
- **iavmd-x.x.x.xxxx-rhel\_9x.src.rpm** – Build the source rpm into the Linux\* VM

§§



## 9 Add Intel® Replacement Driver to Existing VM

When the VM requires a Replacement Driver for bug fix or feature enhancement, the following will apply:

1. Power on VM from ESXi\* Host Client Console.
2. At the GRUB Boot Menu, press 'e' to edit boot parameters. Add to GRUB menu:  
**inst.dd=LABEL=OEMDRV iavmd.direct\_assign=1**  
**initcall\_blacklist=vmd\_drv\_init**

```
Red Hat Enterprise Linux (5.14.0-162.6.1.el9_1.x86_64) 9.1 (Plow)
Red Hat Enterprise Linux (0-rescue-7494563d2a6242daa646c2a468e6c13b) 9.1
UEFI Firmware Settings
```

```
linux ($root)/vmlinuz-5.14.0-162.6.1.el9_1.x86_64 root=/dev/mapper/rhel-root r\
o crashkernel=1G-4G:192M,4G-64G:256M,64G-:512M resume=/dev/mapper/rhel-swap rd\
.lvm.lv=rhel/root rd.lvm.lv=rhel/swap rhgb quiet inst.dd=LABEL=OEMDRV iavmd.di\
rect_assign=1 initcall_blacklist=vmd_drv_init
```

3. Ctrl-x to continue booting VM. Using WinSCP\*, copy the Replacement Driver .rpm file to VM working directory. Install replacement driver.

```
[root@localhost ~]# rpm -ivh kmod-iavmd-1.0.0.1608-rhel_91.x86_64.rpm
```

4. Create a new file on the VM: /etc/modprobe.d/iavmd.conf. Insert **options iavmd direct\_assign=1** in the newly created file. Click Save and exit.

```
options iavmd direct_assign=1
```

5. Update initramfs: **dracut -f**
6. Verify the iavmd file is in initramfs: **lsinitrd /boot/initramfs-\$(uname -r).img |grep iavmd**

```
Last login: Mon Aug 14 15:40:53 2023
[root@localhost ~]# lsinitrd /boot/initramfs-$(uname -r).img | grep iavmd
-rw-r--r-- 1 root root 30 Aug 16 2022 etc/modprobe.d/iavmd.conf
[root@localhost ~]#
```

```
[root@localhost ~]# lsinitrd /boot/initramfs-$(uname -r).img | grep iavmd
-rw-r--r-- 1 root root 44 Aug 16 2022 etc/depmod.d/iavmd-rhel_91.conf
-rw-r--r-- 1 root root 30 Aug 16 2022 etc/modprobe.d/iavmd.conf
-rw-r--r-- 1 root root 6 Aug 16 2022 etc/modules-load.d/iavmd-rhel_91.conf
drwxr-xr-x 2 root root 0 Aug 16 2022 usr/lib/modules/5.14.0-162.6.1.el9_1.x86_64/extra/iavmd
-rw-r--r-- 1 root root 383051 Aug 16 2022 usr/lib/modules/5.14.0-162.6.1.el9_1.x86_64/extra/iavmd/iavmd.ko
```

7. Update the GRUB configuration file. Add the '**initcall\_blacklist=vmd\_drv\_init**' parameter to the end of 'GRUB\_CMDLINE\_LINUX' in **/etc/default/grub** file. This entry may already exist.

```
GRUB_CMDLINE_LINUX="crashkernel=1G-4G:192M,4G-64G:256M,64G-:512M resume=/dev/mapper/rhel-  
swap rd.lvm.lv=rhel/root rd.lvm.lv=rhel/swap rhgb quiet initcall_blacklist=vmd_drv_init"
```

8. Save GRUB configuration: `grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg`

```
[root@localhost ~]# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg  
Generating grub configuration file ...  
Adding boot menu entry for UEFI Firmware Settings ...  
done
```

9. Reboot the VM.
10. Verify Direct Assign devices are present: `nvme list`, `lsblk`, or `ls -la /sys/block`.

§§

## 10 Secure Boot using Intel® Replacement Certificate

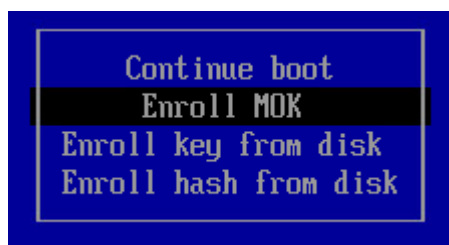
If a replacement driver is used, and Secure Boot is needed, the certificate accompanying the replacement driver must also be installed. Without a valid replacement certificate, DA devices will not display.

### How to install a Replacement Certificate using the Machine Owner Keys (MOK) mokutil tool:


1. Copy the certificate (renamed **intel-vmd-linux-driver.der**) to a VM working directory.
2. Open a Console in the ESXi\* Host Client and power on the target VM.
3. **mokutil --import intel-vmd-linux-driver.der**
4. Enter and confirm a password for the MOK enrollment request. This password will be needed for a later step.

```
[root@localhost Downloads]# mokutil --import intel-vmd-linux-driver.der
input password:
input password again: █
```

5. Reboot the VM.
6. Use the console session from the previous step (not a separate session). Upon reboot, a pending MOK key enrollment request is detected by the shim.efi driver, launches MokManager.efi, and displays the Perform MOK management menu.



7. Select **Enroll MOK** > Continue > Yes to enroll the key, then enter the password created previously when prompted. Select **Reboot**.



8. The public key has now been added to the MOK list. The key is persistent.
9. Power off the VM. Check the box for **Whether or not to enable UEFI secure boot for this VM**.

∨ Boot Options	
Firmware	Choose which firmware should be used to boot the virtual machine: EFI
Enable UEFI secure boot	<input checked="" type="checkbox"/> Whether or not to enable UEFI secure boot for this VM

10. Power on the VM and confirm DA drive visibility.

§§

# 11 Troubleshooting

---

## 11.1 Direct Assigned NVMe's Do Not Display

Power down the VM and Edit Settings. Confirm the PCI device that was selected for Passthrough still displays in Edit settings Virtual Hardware device list.

If it does not display, repeat previous steps on:

1. VM > Virtual Hardware > Add Other Device and select the target DA PCI device.
2. Save the configuration.
3. Reboot VM, confirm PCI device is listed, and recheck for NVMe\* visibility.

## 11.2 VM Does Not Boot with Secure Boot Enabled

Confirm correct security certificate is being used. Inbox RHEL\* VMD driver uses an inbox security certificate. A replacement driver equipped VM requires the accompanying replacement certificate to be imported.

## 11.3 Gen1/2 Xeon System Behavior and Recommendations

Even though this User Guide models a post-Gen1/2 Intel platform DA configuration, some guidance may be useful for older platforms due to non-standard behaviors seen with various ESXi versions. If DA drives do not display as expected, the following guidance may assist:

1. Install iavmd async/replacement driver in Guest OS (ex. kmod-iavmd-1.0.0.1613-rhel\_92.x86\_64 for RHEL 9.2 VM).
2. Add "iavmd.direct\_assign=1" parameter to kernel boot options.
3. If present, remove "8086 201d d3d0 default" controller reference line from VMware host /etc/vmware/passthru.map file and reboot ESXi. Retest for NVMe visibility.

## 11.4 On VM Startup, Error Shows with Failure to Power On due to Wrong Memory Size



1. Ensure Reserve all guest memory (All locked) has been selected in Virtual Hardware > Memory option.
2. Edit the ESXi\* host /vmfs/volumes/[datastore/VM/VM].vmx file to confirm the correct sched.mem.min setting is listed (this setting should be auto-generated to agree with the memory reserved for the VM).

In this case, it should read `sched.mem.min = 16384` since we've assigned 16GB memory to the test VM.

## 11.5 When Using a Replacement Driver in the VM, NVMe\* Devices are Not Visible

Try manually loading the `iaavmd` kernel module: `modprobe -r iaavmd && modprobe iaavmd direct_assign=1`.

## 11.6 Error When Installing RHEL OS to VM

If a VMware\* error displays when installing RHEL OS to the VM, isolate by removing any DA controller/PCI Devices from Host HW. Ensure datastore is not located on a DA disk. Install RHEL on standard (non-DA) VM, then add DA PCI Devices one at a time to detect where the issue may be.

§§