

How-to: Configure NVIDIA GPU device in VMDirectPath I/O passthrough mode on VMware ESXi 6.x server.

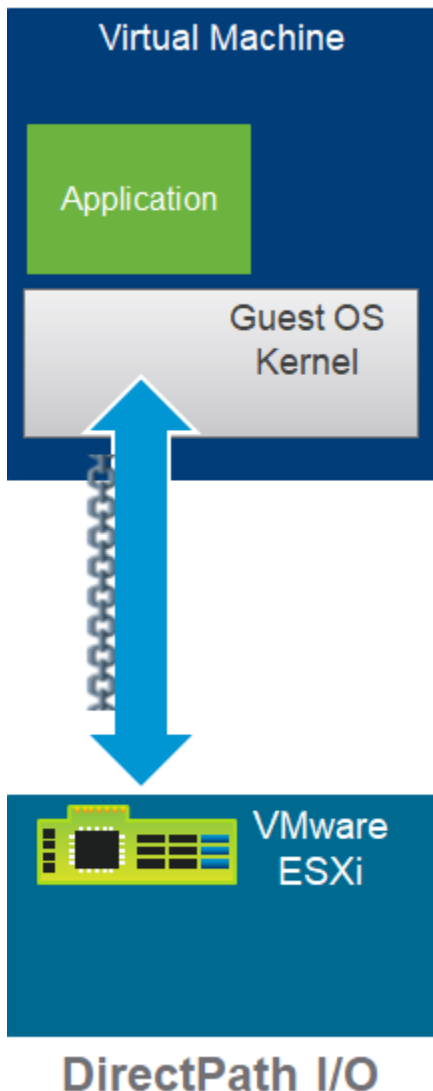
Created on Jun 9, 2019

Introduction

This post describes the procedure of how to configure Nvidia GPU device VMDirectPath I/O passthrough mode on VMware ESXi 6.x server.

About VMDirectPath I/O passthrough devices

16 Devices is supported maximum for passthrough devices per VMs on ESXi 6.x.



Related Documents


- [How-to: Install NVIDIA Firmware Tools \(MFT\) on VMware ESXi 6.7 /7.0.](#)
- [RDG: RoCE accelerated vSphere 6.7 cluster deployment for ML and HPC workloads.](#)
- [RDG: VMware NSX-V hardware VTEPs in High-Availability mode on Spectrum switches running Cumulus Linux.](#)
- [How-to: Firmware update for NVIDIA ConnectX-5/6 adapter on VMware ESXi 6.5 and above.](#)
- [How-to: Change port type of NVIDIA ConnectX VPI adapter on VMware ESXi 6.x and above.](#)
- [How-to: NVIDIA ConnectX driver upgrade on VMware ESXi 6.7/7.0 and above.](#)
- [How-to: Configure NVIDIA network device in VMDirectPath I/O passthrough mode on VMware ESXi 6.x.](#)
- [How-to: Configure NVIDIA ConnectX-5/6 adapter in SR-IOV mode on VMware ESXi 6.7/7.0 and above.](#)
- [How-to: Configure RoCEv2 lossless fabric for VMware ESXi 6.5 and above.](#)
- [How-to: Configure NVIDIA GPU device in VMDirectPath I/O passthrough mode on VMware ESXi 6.x server.](#)
- [How-to: Configure PVRDMA in VMware vSphere 6.5/6.7.](#)
- [RDG: Kubernetes Cluster Deployment for ML and HPC Workloads with NVIDIA GPU Virtualization and VMware PVRDMA Technologies.](#)
- [How-to: Configure NVIDIA network device in DirectPath I/O and Dynamic DirectPath I/O passthrough modes on VMware ESXi 7.0.](#)
- [How-to: Configure RoCE PVRDMA Namespace in VMware vSphere 7.0.](#)
- [RDG: Configure NVMe-oF RoCE Datastore in VMware vSphere 7.0 with Pavilion Storage.](#)

Configuring passthrough Nvidia GPU device

Before configuring a device for PCI Passthrough, ensure that platform and device meets the PCI Passthrough requirements, see [VMware vSphere VMDirectPath I/O: Requirements for Platforms and Devices \(2142307\)](#).

To configure passthrough devices on an ESXi host:

1. **Select an ESXi host** from the Inventory of VMware vSphere Client.
2. **Enter Maintenance Mode** the ESXi host.
3. In the **Navigator** click **Manage, Hardware** and **PCI Devices**. The passthrough Configuration page lists all available pass-through devices.
4. **Select the devices** and click **Toggle passthrough**.



5. **Reboot** host for the change to take effect. After rebooting, the devices are marked with a green icon and are enabled.



Note: The configuration changes are saved in the `/etc/vmware/esx.conf` file.


The PCI slot number where the device was connected is `04:00.0`. It is recorded as:

```
/device/00000:004:00.0/owner == "passthru"
```

6. To **verify** it:

In UI:

- a. **Select an ESXi host** from the Inventory of VMware vSphere Client.
- b. **Enter Maintenance Mode** the ESXi host.
- c. In the **Navigator** click **Manage, Hardware** and **PCI Devices**. The passthrough Configuration page lists all available pass-through devices.
- d. **Verify:**



In CLI:

- a. **Enable SSH Access** to ESXi server.
- b. **Log into ESXi vSphere Command-Line Interface** with root permissions.
- c. **Edit the `/etc/vmware/esx.conf`** file by changing the owner of the PCI device from passthru to vmkernel.

ESXi Console

```
# vi /etc/vmware/esx.conf
```

Find `/device/<PCI_ID>/owner = "passthru"`.

Write:
`/passthru`

- d. **Check:**

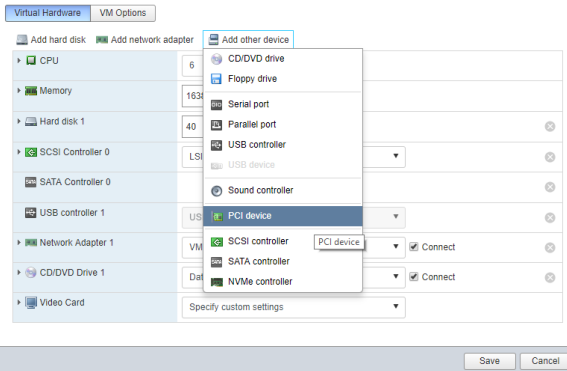
/device/00000:004:00.0/owner == "passthru"

6. Reboot the ESXi host.

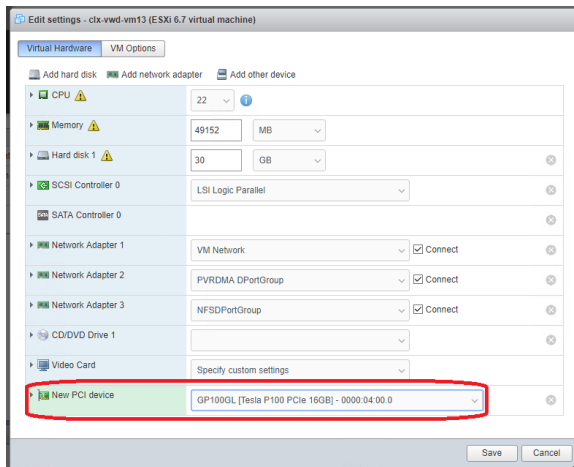
7. Exit Maintenance Mode the ESXi host.

To configure a PCI device on a virtual machine:

1. From the **Navigator** in vSphere Client, **right-click** the **virtual machine** and click **Edit Settings**.
2. Select the **Virtual Hardware** tab.
3. Click **Add other device**.



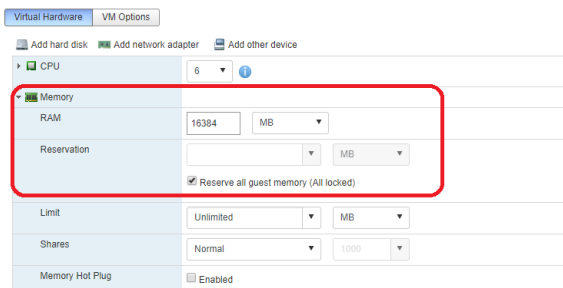
Choose the **PCI Device**.



Click **Save**.



Note: When the device is assigned, the virtual machine must have a memory reservation for the full configured memory size.



Done!

Configuring non-passthrough Nvidia GPU device

To force the change to non-passthrough mode, complete these steps:

1. **Enable SSH Access to ESXi** server.
2. **Log into ESXi vSphere Command-Line Interface** with root permissions.
3. **Enter Maintenance Mode** the ESXi host.
4. Edit the `/etc/vmware/esx.conf` file by changing the owner of the PCI device from passthru to vmkernel.

ESXi Console

```
# vi /etc/vmware/esx.conf
```

5. **Find** `/device/<PCI_ID>/owner = "passthru"`.
6. **Change from:**

ESXi Console

```
/device/00000:004:00.0/owner = = "passthru"
```

Change to:

ESXi Console

```
/device/00000:004:00.0/owner = = "vmkernel"
```


7. **Reboot** the ESXi host.

8. **Verify.**

ESXi Console

```
# lspci | grep Mellanox
```

```
0000:04:00.0 Display controller: NVIDIA Corporation GP100GL [Tesla P100 PCIe 16GB]
```



9. **Exit Maintenance Mode** the ESXi host.

Done!