

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

Created on Jun 9, 2019

Introduction

This post describes the procedure of how to configure NVIDIA network device into and from VMDirectPath I/O passthrough mode on VMware ESXi 6.x.

References

- [NVIDIA ConnectX-4/ConnectX-5 NATIVE ESXi Driver for VMware vSphere 6.5 User Manual](#)
- [NVIDIA ConnectX-4/ConnectX-5 NATIVE ESXi Driver for VMware vSphere 6.7 User Manual](#)
- [ConnectX® Ethernet Driver for VMware® ESXi Server](#)
- [vSphere Command-Line Interface Concepts and Examples](#)

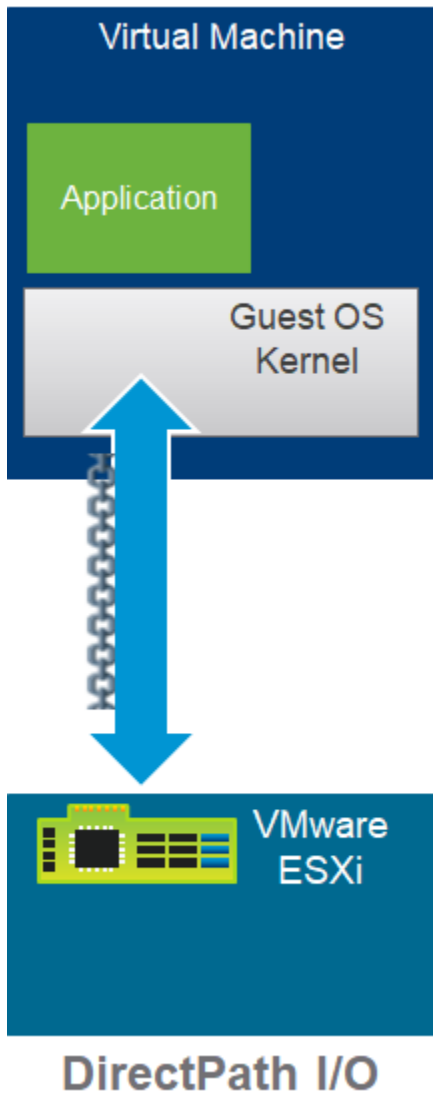
About VMDirectPath I/O passthrough devices

The supported **maximums** for passthrough devices per VMs are:

- ESXi **6.5**: **16** Devices
- ESXi **6.7**: **16** Devices

Related Documents

- [How-to: Install NVIDIA Firmware Tools \(MFT\) on VMware ESXi 6.7 /7.0.](#)
- [RDG: RoCE accelerated Apache Spark 2.2 cluster deployment.](#)
- [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.](#)
- [QSG: TRex in a few steps using Nvidia ConnectX adapters.](#)
- [QSG: Building Docker image with compiled Nvidia Network DPDK PMD.](#)
- [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.](#)




Configuring passthrough NVIDIA network 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 ESX/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 **02:00.1**. It is recorded as:

```
/device/00000:002:00.1/owner = = "passthru"
```

6. To **verify** in:

UI:

- Select an **ESXiESXi** host from the Inventory of VMware vSphere Client.
- In the **Navigator** click **Manage, Hardware and PCI Devices**. The passthrough Configuration page lists all available pass-through devices.
- Verify:**



In CLI:

- Enable **SSH Access to ESXi** server.
- Log into **ESXi vSphere Command-Line Interface** with root permissions.
- View 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"`.

d. **Check that:**

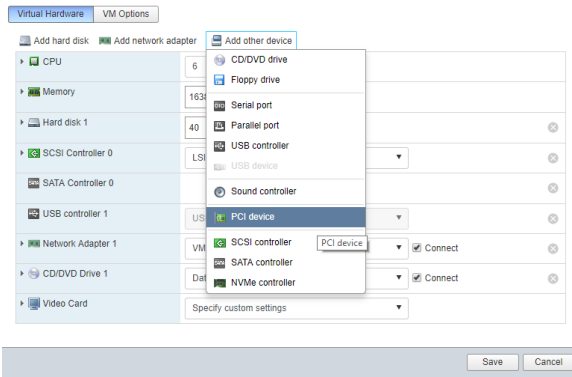
ESXi Console

```
/device/00000:002:00.1/owner = = "passthru"
```

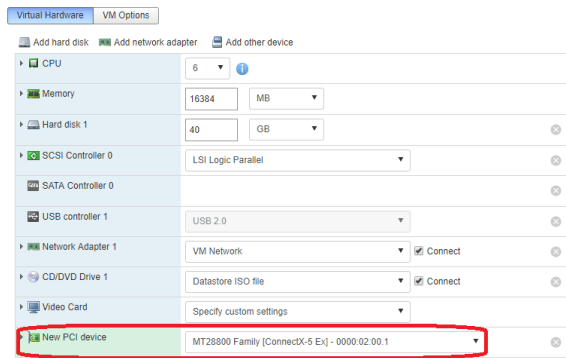
- Reboot** the ESXi host.
- Exit Maintenance Mode** the ESXi host.

To configure a PCI device on a virtual machine:

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

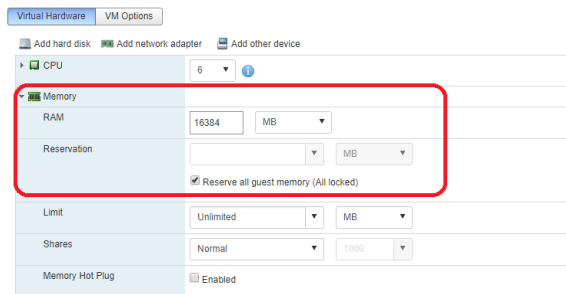


Choose the PCI Device.



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 network 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
```

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

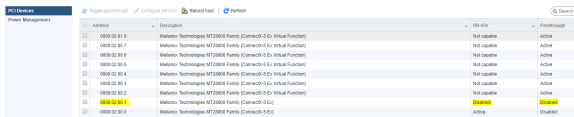
ESXi Console

```
/passthru
```

- Change from:
`/device/0000:002:00.1/owner == "passthru"`
Change to:
`/device/0000:002:00.1/owner == "vmkernel"`
- Reboot the ESXi host.
- Verify.

ESXi Console

```
# lspci | grep Mellanox  
  
0000:02:00.0 Network controller: Mellanox Technologies MT28800 Family  
[ConnectX-5 Ex] [vmnic2]  
0000:02:00.1 Network controller: Mellanox Technologies MT28800 Family  
[ConnectX-5 Ex] [vmnic3]
```



Address	Description	VM NIC	Provisioning
0000:02:00.0	Mellanox Technologies MT28800 Family ConnectX-5 Ex Virtual Function	Not available	Active
0000:02:00.1	Mellanox Technologies MT28800 Family ConnectX-5 Ex Virtual Function	Not available	Active
0000:02:00.2	Mellanox Technologies MT28800 Family ConnectX-5 Ex Virtual Function	Not available	Active
0000:02:00.3	Mellanox Technologies MT28800 Family ConnectX-5 Ex Virtual Function	Not available	Active
0000:02:00.4	Mellanox Technologies MT28800 Family ConnectX-5 Ex Virtual Function	Not available	Active
0000:02:00.5	Mellanox Technologies MT28800 Family ConnectX-5 Ex Virtual Function	Not available	Active
0000:02:00.6	Mellanox Technologies MT28800 Family ConnectX-5 Ex Virtual Function	Not available	Active
0000:02:00.7	Mellanox Technologies MT28800 Family ConnectX-5 Ex Virtual Function	Not available	Active
0000:02:00.8	Mellanox Technologies MT28800 Family ConnectX-5 Ex Virtual Function	Not available	Active
0000:02:00.9	Mellanox Technologies MT28800 Family ConnectX-5 Ex	Active	Disabled

- Exit Maintenance Mode the ESXi host.

Done!