

proxmox

gpu

passthrough

Authors: [fire1ce](#) | Created: [2021-08-27](#) | Last update: [2023-07-13](#)

Proxmox GPU Passthrough to VM

Introduction

GPU passthrough is a technology that allows the Linux kernel to present the internal PCI GPU directly to the virtual machine. The device behaves as if it were powered directly by the virtual machine, and the virtual machine detects the PCI device as if it were physically connected. We will cover how to enable GPU passthrough to a virtual machine in Proxmox VE.

Your mileage may vary depending on your hardware.

Proxmox Configuration for GPU Passthrough

The following examples uses `SSH` connection to the Proxmox server. The editor is `nano` but feel free to use any other editor. We will be editing the `grub` configuration file.

Find the PCI address of the GPU Device. The following command will show the PCI address of the GPU devices in Proxmox server:

```
lspci -nnv | grep VGA
```

Find the GPU you want to passthrough in result ts should be similar to this:

```
01:00.0 VGA compatible controller [0300]: NVIDIA Corporation TU104 [GeForce RTX 2080 SUPER] [10de:1e81] (rev a1) (prog-if 00 [VGA controller])
```

What we are looking is the PCI address of the GPU device. In this case it's `01:00.0`.

`01:00.0` is only a part of a group of PCI devices on the GPU.

We can list all the devices in the group `01:00` by using the following command:

```
lspci -s 01:00
```

The usual output will include VGA Device and Audio Device. In my case, we have a USB

Controller and a Serial Bus controller.

Get 50% off Torguard VPN With Coupon: [all50torgourd](#)



```
01:00.0 VGA compatible controller: NVIDIA Corporation TU104 [GeForce RTX 2080 SUPER] (rev a1)
01:00.1 Audio device: NVIDIA Corporation TU104 HD Audio Controller (rev a1)
01:00.2 USB controller: NVIDIA Corporation TU104 USB 3.1 Host Controller (rev a1)
01:00.3 Serial bus controller [0c80]: NVIDIA Corporation TU104 USB Type-C UCSI Controller (rev a1)
```

Now we need to get the id's of those devices. We can do this by using the following command:

```
lspci -s 01:00 -n
```

The output should look similar to this:

```
01:00.0 0300: 10de:1e81 (rev a1)
01:00.1 0403: 10de:10f8 (rev a1)
01:00.2 0c03: 10de:1ad8 (rev a1)
01:00.3 0c80: 10de:1ad9 (rev a1)
```

What we are looking are the pairs, we will use those id to split the PCI Group to separate devices.

```
10de:1e81,10de:10f8,10de:1ad8,10de:1ad9
```

Now it's time to edit the `grub` configuration file.

```
nano /etc/default/grub
```

Find the line that starts with `GRUB_CMDLINE_LINUX_DEFAULT` by default they should look like this:

```
GRUB_CMDLINE_LINUX_DEFAULT="quiet"
```

For Intel CPU

```
intel_iommu=on
```

For AMD CPU

```
amd_iommu=on
```

Then change it to look like this (Intel CPU example) and replace `vfio-pci.ids=` with the ids for the GPU you want to passthrough:

```
vfio-pci.ids=10de:1e81,10de:10f8,10de:1ad8,10de:1ad9
vfio_iommu_type1.allow_unsafe_interrupts=1 kvm.ignore_msrs=1
modprobe.blacklist=radeon,nouveau,nvidia,nvidiafb,nvidia-gpu"
```

Save the config changed and then update GRUB.

```
update-grub
```

Next we need to add `vfio` modules to allow PCI passthrough.

Edit the `/etc/modules` file.

```
nano /etc/modules
```

Add the following line to the end of the file:

```
# Modules required for PCI passthrough
vfio
vfio_iommu_type1
vfio_pci
vfio_virqfd
```

Save and exit the editor.

Update configuration changes made in your `/etc` filesystem

```
update-initramfs -u -k all
```

Reboot Proxmox to apply the changes

Verify that IOMMU is enabled

```
dmesg | grep -e DMAR -e IOMMU
```

There should be a line that looks like `DMAR: IOMMU enabled`. If there is no output, something is wrong.

```
[0.000000] Warning: PCIe ACS overrides enabled; This may allow non-IOMMU
protected peer-to-peer DMA
[0.067203] DMAR: IOMMU enabled
[2.573920] pci 0000:00:00.2: AMD-Vi: IOMMU performance counters supported
[2.580393] pci 0000:00:00.2: AMD-Vi: Found IOMMU cap 0x40
[2.581776] perf/amd_iommu: Detected AMD IOMMU #0 (2 banks, 4 counters/bank).
```

Check that the GPU is in a separate IOMMU Group by using the following command:

```
#!/bin/bash
```

 Get 50% off TorGuard VPN With Coupon: [all50torgourd](#)



```
for g in $(find /sys/kernel/iommu_groups/* -maxdepth 0 -type d | sort -V); do
  echo "IOMMU Group ${g##*/}:"
  for d in $g/devices/*; do
    echo -e "\t$(lspci -nns ${d##*/})"
  done;
done;
```

Now your Proxmox host should be ready to GPU passthrough!

Windows Virtual Machine GPU Passthrough Configuration

For better results its recommend to use this [Windwos 10/11 Virutal Machine configuration for proxmox.](#)

✖ Limitations & Workarounds

- In order for the GPU to to function properly in the VM, you must disable Proxmox's Virutal Display - Set it `none` .
- You will lose the ability to conect to the VM via Proxmox's Console.
- Display must be conected to the physical output of the GPU for the Windows Host to initialize the GPU properly.
- **You can use a [HDMI Dummy Plug](#) as a workaround - It will present itself as a HDMI Display to the Windows Host.**
- Make sure you have alternative way to connect to the VM for example via Remote Desktop (RDP).

Find the PCI address of the GPU.

```
lspci -nnv | grep VGA
```

This should result in output similar to this:

```
01:00.0 VGA compatible controller [0300]: NVIDIA Corporation TU104 [GeForce RTX 2080 SUPER] [10de:1e81] (rev a1) (prog-if 00 [VGA controller])
```

If you have multiple VGA, look for the one that has the `Intel` in the name.

Here, the PCI address of the GPU is `01:00.0` .

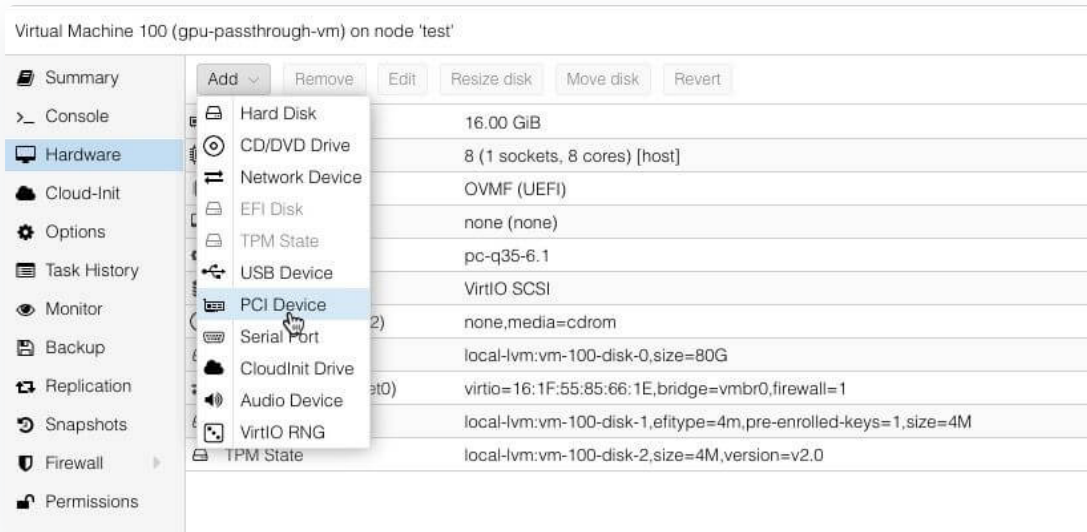
```
root@pve:~# lspci -nnv | grep VGA
00:02.0 VGA compatible controller [0300]: Intel Corporation CometLake-S GT2 [UHD Graphics 630] [8086:3e92] (prog-if 00 [VGA controller])
01:00.0 VGA compatible controller [0300]: NVIDIA Corporation TU104 [GeForce RTX 2080 SUPER] [10de:1e81] (rev a1) (prog-if 00 [VGA controller])
root@pve:~#
```

For best performance the VM should be configured the `Machine` type to `q35`.

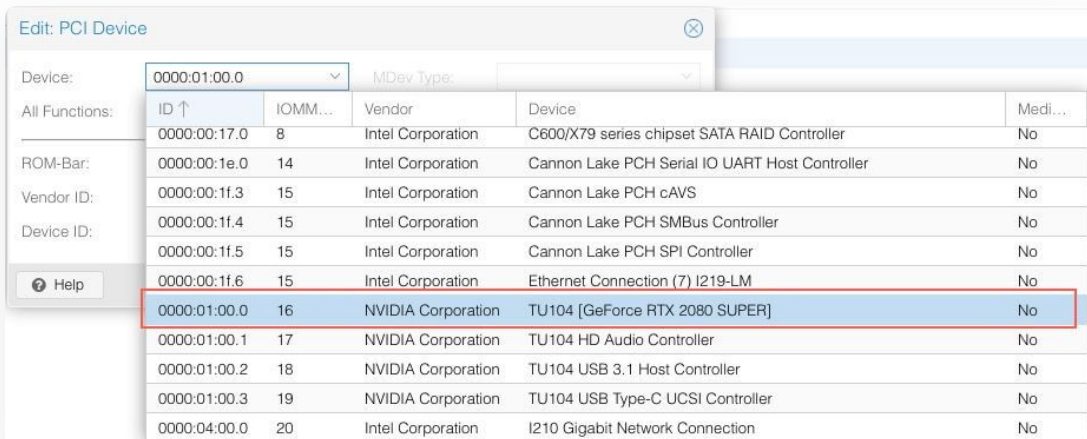
🗨️ Get 50% off Torguard VPN With Coupon: [all50torgourd](#)



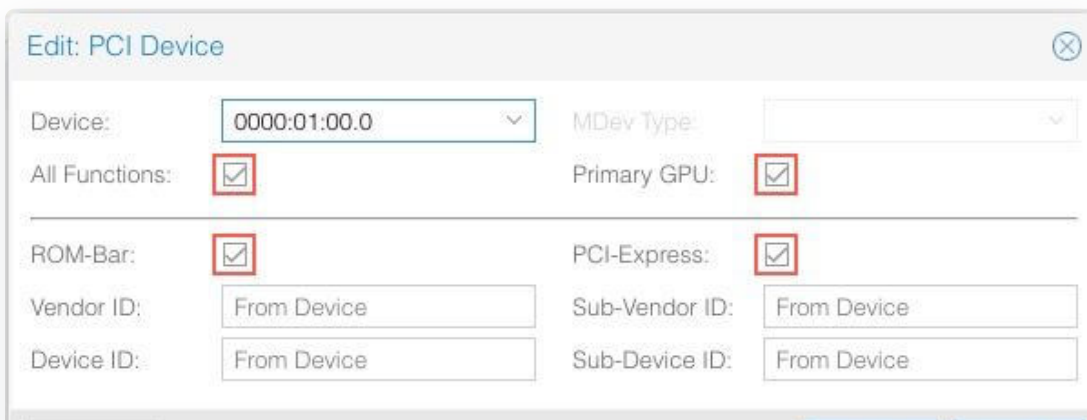
Open the web gui and navigate to the **Hardware** tab of the VM you want to add a vGPU. Click **Add** above the device list and then choose **PCI Device**



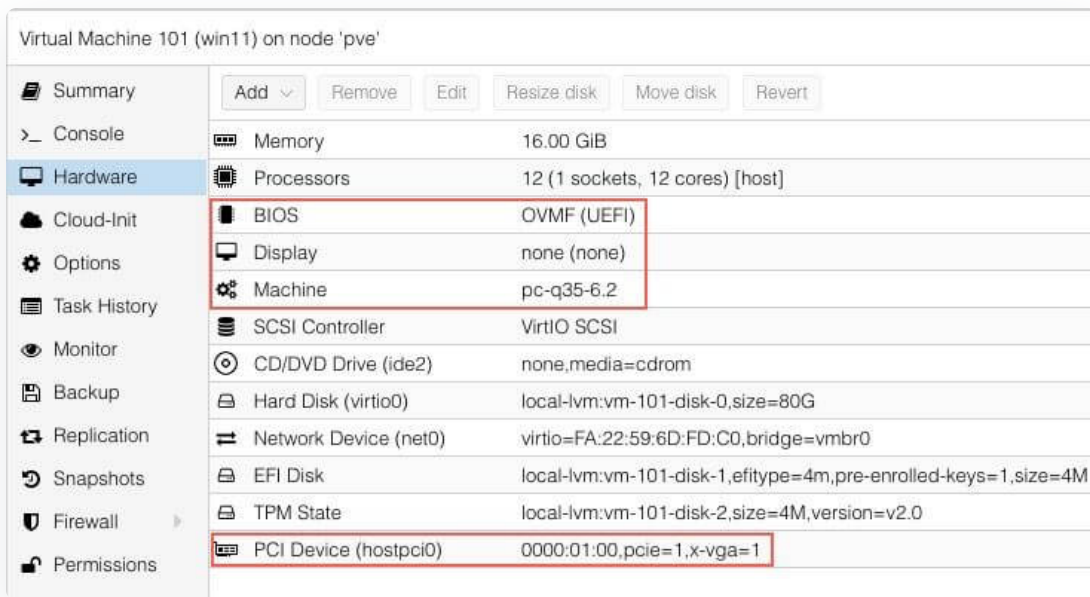
Open the **Device** dropdown and select the GPU, which you can find using it's PCI address. This list uses a different format for the PCI addresses id, `01:00.0` is listed as `0000:01:00.0`.



Select **All Functions**, **ROM-Bar**, **Primary GPU**, **PCI-Express** and then click **Add**.




The Windows Virtual Machine Proxmox Setting should look like this:



Power on the Windows Virtual Machine.

Connect to the VM via Remote Desktop (RDP) or any other remote access protocol you prefer. Install the latest version of GPU Driver for your GPU.

If all when well you should see the following output in `Device Manager` and `GPU-Z`:



That's it!

Linux Virtual Machine GPU Passthrough Configuration

We will be using Ubuntu Server 20.04 LTS. for this guide.

Get 50% off Torguard VPN With Coupon: [all50torguard](#)



```
lspci -nnv | grep VGA
```

This should result in output similar to this:

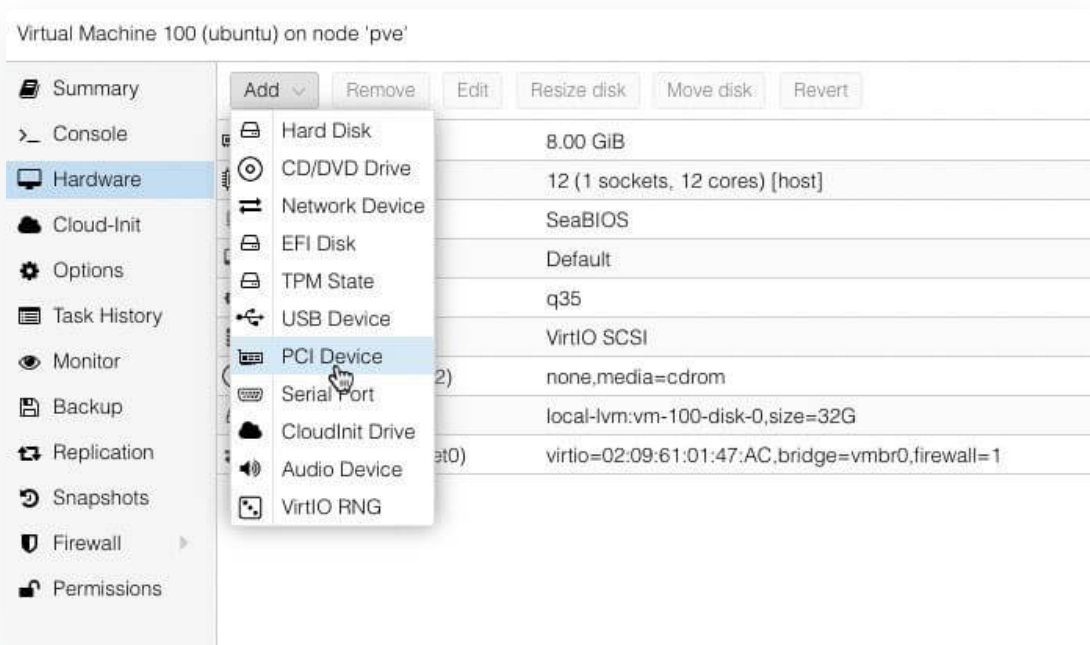
```
01:00.0 VGA compatible controller [0300]: NVIDIA Corporation TU104 [GeForce RTX 2080 SUPER] [10de:1e81] (rev a1) (prog-if 00 [VGA controller])
```

If you have multiple VGA, look for the one that has the `Intel` in the name. Here, the PCI address of the GPU is `01:00.0`.

```
root@pve:~# lspci -nnv | grep VGA
00:02.0 VGA compatible controller [0300]: Intel Corporation CometLake-S GT2 [UHD Graphics 630] [8086:3e92] (prog-if 00 [VGA controller])
01:00.0 VGA compatible controller [0300]: NVIDIA Corporation TU104 [GeForce RTX 2080 SUPER] [10de:1e81] (rev a1) (prog-if 00 [VGA controller])
root@pve:~#
```

For best performance the VM should be configured the `Machine` type to `q35`.

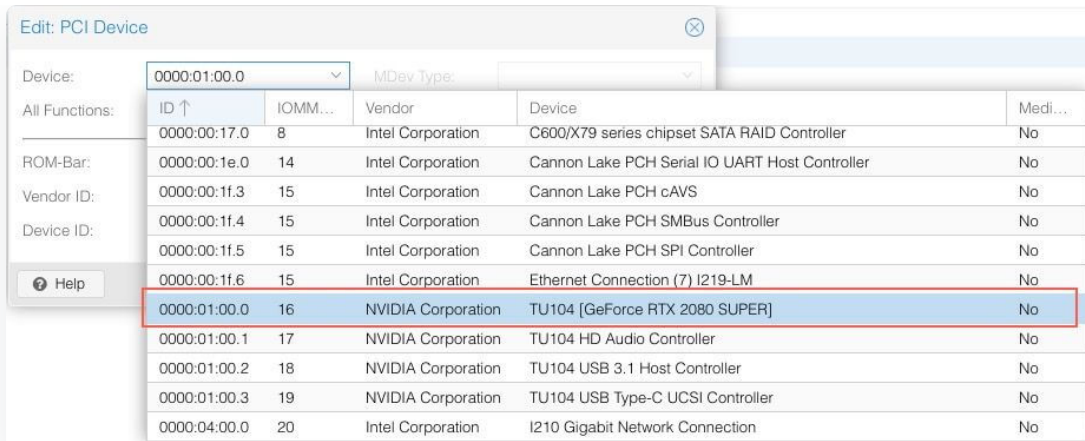
This will allow the VM to utilize PCI-Express passthrough.



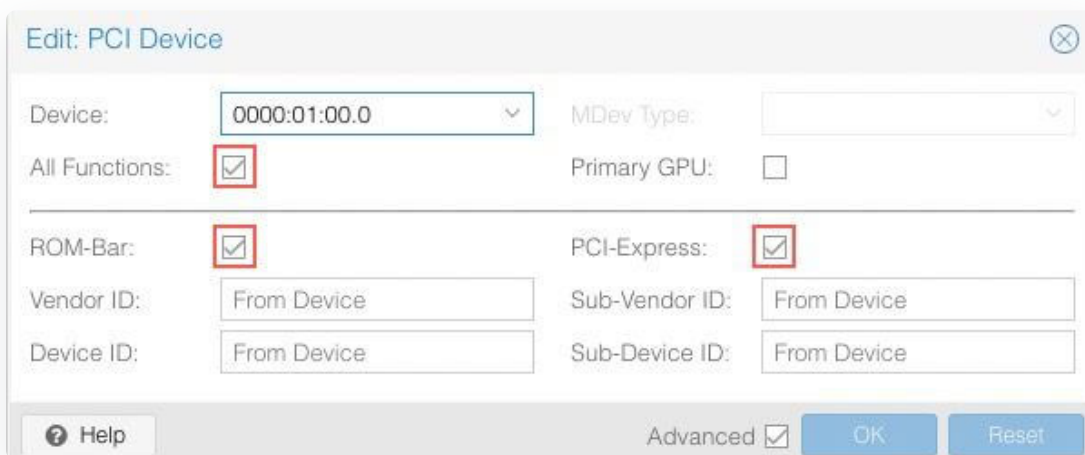
Open the `Device` dropdown and select the GPU, which you can find using it's PCI address.

This list uses a different format for the PCI addresses id, `01:00.0` is listed as

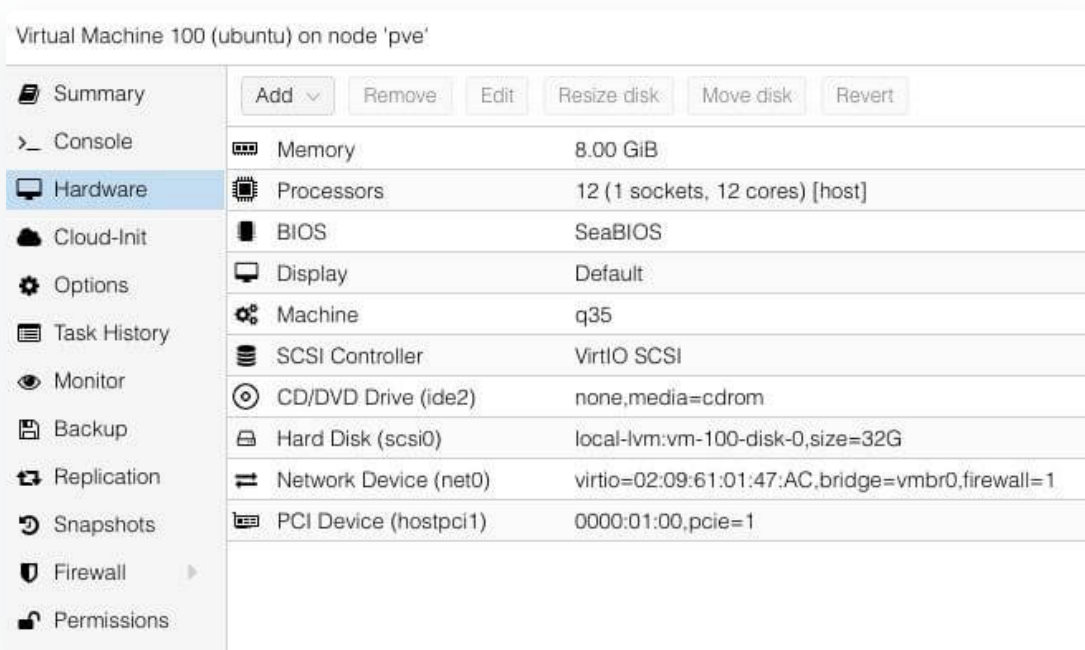
```
0000:01:00.0
```

Select All Functions , ROM-Bar , PCI-Epress and then click Add .



The Ubuntu Virtual Machine Proxmox Setting should look like this:



Boot the VM. To test the GPU passthrough was successful, you can use the following

Get 50% off Torguard VPN With Coupon: [all50torgourd](#) ✕


```
sudo lspci -nnv | grep VGA
```

The output should include the GPU:

```
01:00.0 VGA compatible controller [0300]: NVIDIA Corporation TU104 [GeForce  
RTX 2080 SUPER] [10de:1e81] (rev a1) (prog-if 00 [VGA controller])
```

Now we need to install the GPU Driver. I'll be covering the installation of Nvidia Drivers in the next example.

Search for the latest Nvidia Driver for your GPU.

```
sudo apt search nvidia-driver
```

In the next step we will install the Nvidia Driver v535.

Note

--no-install-recommends is important for Headless Server. `nvidia-driver-535` will install `xorg` (GUI) `--no-install-recommends` flag will prevent the GUI from being installed.

```
sudo apt install --no-install-recommends -y build-essential nvidia-driver-535  
nvidia-headless-535 nvidia-utils-535 nvidia-cuda-toolkit
```

This will take a while to install. After the installation is complete, you should reboot the VM.

Now let's test the Driver initialization. Run the following command in the VM:

```
nvidia-smi && nvidia-smi -L
```

If all went well you should see the following output:

```

fire1ce@ibuntu:~$ nvidia-smi && nvidia-smi -L
Sun Apr 17 11:21:36 2022
+-----+
| NVIDIA-SMI 510.54      Driver Version: 510.54      CUDA Version: 11.6      |
+-----+-----+-----+
| GPU  Name                Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
|                                           |                  |     MIG M. |
+-----+-----+-----+
|   0   NVIDIA GeForce ...   Off   | 00000000:01:00.0 Off  |           N/A   |
|  0%   66C    P0     36W / 250W |  0MiB /  8192MiB |    0%    Default  |
|                                           |                  |           N/A   |
+-----+-----+-----+

Processes:
+-----+-----+-----+
| GPU  GI  CI           PID  Type  Process name          GPU Memory |
|   ID  ID  ID                   |                |           Usage |
+-----+-----+-----+
| No running processes found |
+-----+-----+-----+
GPU 0: NVIDIA GeForce RTX 2080 SUPER (UUID: GPU-1925cf42-3927-b6b3-e3f6-a76316680682)
fire1ce@ibuntu:~$

```

That's it! You should now be able to use the GPU for hardware acceleration inside the VM.

Debug

Dbug Messages - Shows Hardware initialization and errors

```
dmesg -w
```

Display PCI devices information

```
lspci
```

Display Driver in use for PCI devices

```
lspci -k
```

Display IOMMU Groups the PCI devices are assigned to

```

#!/bin/bash
shopt -s nullglob
for g in $(find /sys/kernel/iommu_groups/* -maxdepth 0 -type d | sort -V); do
    echo "IOMMU Group ${g##*/}:"
    for d in $g/devices/*; do
        echo -e "\t${lspci -nns ${d##*/}}"
    done;
done;

```

Reboot Proxmox to apply the changes

Get 50% off TorGuard VPN With Coupon: [all50torgourd](#)



Comments

0 reactions




1 comment – powered by giscus

Oldest

Newest



 Get 50% off Torguard VPN With Coupon: [all50torgourd](#) 