

Proxmox GPU Passthrough

GPU Passthrough

This configuration worked for me, you might need to change things around

Keep in mind I have an AMD CPU and Nvidia GPU, if you have other config, you might have to use different commands

After upgrading Proxmox to 7.2, passthrough wasn't working. To make it work again try resetting your graphics card: [Resetting GPU](#)
OR keep reading, the GRUB parameters have to be changed to make it work again with the latest kernel!

Configuring BIOS

Before doing anything make sure virtualization and IOMMU is enabled in your BIOS, you can't do anything before that.

If your motherboard doesn't support IOMMU, then you can't pass through PCI(e) devices to your VMs.

Update the Host configuration

Login to the host and open `/etc/default/grub`. Find the line `GRUB_CMDLINE_LINUX_DEFAULT` and change it from:

```
GRUB_CMDLINE_LINUX_DEFAULT="quiet"
```

to

```
GRUB_CMDLINE_LINUX_DEFAULT="quiet iommu=pt nofb nomodeset video=vesafb:off video=efifb:off"
```

Please note that `vesafb:off` and `efifb:off` are assigned to `video` in two different statements. Many tutorials mention a shorted version, like `video=vesafb:off,efifb:off` but that didn't work for me.

Rather than using "video=vesafb:off video=efifb:off" parameters, try replacing them with "initcall_blacklist=sysfb_init" when using the latest kernel (5.15.64-1-pve), so the GPU passthrough will work again

Run `update-grub` to append the grub's content to all linux entries in `/boot/grub/grub.cfg`.

Next open `/etc/modules` and add the followings:

```
vfi  
vfi_iommu_type1  
vfi_pci  
vfi_virqfd
```

After these changes run the below to refresh the `initramfs`, then restart your server:

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

Once it's restarted, run the below commands to check if IOMMU was successfully enabled:

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

It should display that `IOMMU, Directed I/O or Interrupt Remapping is enabled` or something similar, it could be different on your hardware.

Also check that the devices are in different IOMMU groups:

```
find /sys/kernel/iommu_groups/ -type l
```

Device passthrough setup

First find the device Ids that you want to passthrough.

Run

```
lspci -nn
```

which will display all the devices and their Ids in the host. Find yours and write it down.

It looks something like `[1245:4f5a]`, don't forget to copy the audio device's Id as well

Since we want to use a GPU in our VM, we have to passthrough both the video and sound devices, you can't pass only one of them.

You also have to blacklist your GPU so the host won't utilize it. This is how my `/etc/modprobe.d/pve-blacklist.conf` looks like:

```
# This file contains a list of modules which are not supported by Proxmox VE

# nvidiafb see bugreport https://bugzilla.proxmox.com/show_bug.cgi?id=701
blacklist nvidiafb

#blacklist radeon # Use this line if you have an AMD card
#blacklist amdgpu # Use this line if you have an AMD card
blacklist nvidia
blacklist nouveau
```

Then in your `/etc/modprobe.d/vfio.conf` insert:

```
options vfio-pci ids=1245:4f5a,1002:aad8 disable_vga=1
```

Here the two ids are the ones which you copied previously.

Also create `/etc/modprobe.d/kvm.conf` with the below content:














```
options kvm ignore_msrs=1
```

This will allow to use Nvidia cards on Windows when you set the CPU to host.

Apply these changes: `update-initramfs -u -k all` then restart the host.

At this point your host should be ready.

Creating VM

 Memory	12.00 GiB
 Processors	4 (1 sockets, 4 cores) [kvm64,flags=+aes] [cpuunits=3096]
 BIOS	OVMF (UEFI)
 Display	none (none)
 Machine	q35
 SCSI Controller	VirtIO SCSI
 CD/DVD Drive (ide2)	none,media=cdrom
 Hard Disk (scsi0)	container:136/vm-136-disk-0.qcow2,size=52G
 Network Device (net0)	virtio=A2:0B:1E:90:B5:B8,bridge=vbr0,firewall=1
 EFI Disk	container:136/vm-136-disk-1.qcow2,size=128K
 USB Device (usb0)	host=046d:c05b,usb3=1
 USB Device (usb1)	host=413c:2113,usb3=1
 PCI Device (hostpci0)	07:00,pcie=1,x-vga=1

Configuration in a text format:

```
bios: ovmf
bootdisk: scsi0
cores: 4
cpu: kvm64,flags=+aes
cpuunits: 3096
efidisk0: container:136/vm-136-disk-1.qcow2,size=128K
hostpci0: 07:00,pcie=1,x-vga=1
ide2: none,media=cdrom
machine: q35
memory: 12288
name: ubuntu8
net0: virtio=A2:0B:1E:90:B5:B8,bridge=vbr0,firewall=1
numa: 0
ostype: l26
scsi0: container:136/vm-136-disk-0.qcow2,size=52G
scsihw: virtio-scsi-pci
smbios1: uuid=b925668d-9785-4941-ab36-4151164248c7
sockets: 1
usb0: host=046d:c05b,usb3=1
usb1: host=413c:2113,usb3=1
vga: none
vmgenid: 9aae2c4f-30ff-4a2b-ac56-805e49c670d5
```

- The `07:00` is my GPU set to `hostpci0`
- vga has to be set to null
- cpu can be kvm64, it doesn't have to be host
- I set cpuunits to higher than default so proxmox will prioritize this VM

Sources:

- https://pve.proxmox.com/wiki/Pci_passthrough
- [https://pve.proxmox.com/wiki/PCI\(e\)_Passthrough](https://pve.proxmox.com/wiki/PCI(e)_Passthrough)
- https://old.reddit.com/r/homelab/comments/b5xpua/the_ultimate_beginners_guide_to_gpu_passthrough/
- https://wiki.archlinux.org/index.php/PCI_passthrough_via_OVMF
- <https://www.kernel.org/doc/html/v5.15/admin-guide/kernel-parameters.html>
- <https://forum.proxmox.com/threads/problem-with-gpu-passthrough.55918/post-478351>

Revision #7

Created 30 April 2022 18:22:23 by Tozo

Updated 29 October 2022 16:34:07 by Tozo