Nvidia GPU Pass-through to Container in Docker VM

1.) Install driver in VM

2 Fehler bei Treiberinstallation beheben:

- blacklist nouveau Treiber

```
# Datei erstellen
nano /etc/modprobe.d/blacklist-nouveau.conf
#Inhalt der Datei
blacklist nouveau
options nouveau modeset=0
# Update initframs
sudo update-initramfs -u
# Neu starten
reboot
```

- gcc-compiler installieren

```
# Install dependencies
sudo apt-get install build-essential gcc-multilib dkms
```

Kompatible Treiber-Versionen + Unlock-Patch installieren: https://github.com/keylase/nvidia-patch

```
#Directory erstellen
mkdir /opt/nvidia && cd /opt/nvidia
#Treiber download
wget https://international.download.nvidia.com/XFree86/Linux-x86_64/525.89.02/NVIDIA-Linux-x86_64-525.89.02.run
#Mount
chmod +x ./NVIDIA-Linux-x86_64-525.89.02.run
#Execute
./NVIDIA-Linux-x86_64-525.89.02.run
```

2.) Setting up NVIDIA Container Toolkit in VM

Setup the package repository and the GPG key:

```
$ distribution=$(. /etc/os-release;echo $ID$VERSION_ID) \
      && curl -fsSL https://nvidia.github.io/libnvidia-container/gpgkey | sudo gpg --dearmor -o /usr/share/keyrings/nvidia-container-toolkit-keyring.gpg \
      && curl -s -L https://nvidia.github.io/libnvidia-container/$distribution/libnvidia-container.list | \
            sed 's#deb https://#deb [signed-by=/usr/share/keyrings/nvidia-container-toolkit-keyring.gpg] https://#g' | \
            sudo tee /etc/apt/sources.list.d/nvidia-container-toolkit.list
```

Install `nvidia-container-toolkit` package (and dependencies):

```
# Install the nvidia-container-toolkit package (and dependencies) after updating the package listing:
$ sudo apt-get update
$ sudo apt-get install -y nvidia-container-toolkit

# Configure the Docker daemon to recognize the NVIDIA Container Runtime:
$ sudo nvidia-ctk runtime configure --runtime=docker

# Restart the Docker daemon to complete the installation after setting the default runtime:
$ sudo systemctl restart docker

# At this point, a working setup can be tested by running a base CUDA container:
$ sudo docker run --rm --runtime=nvidia --gpus all nvidia/cuda:11.6.2-base-ubuntu20.04 nvidia-smi
```

3.) Nvidia-GPU pass-through in VM

docker-compose erweitern mit "runtime: nvidia" und "NVIDIA_VISIBLE_DEVICES=all":

```
version: "2.1"
services:
  jellyfin:
    image: lscr.io/linuxserver/jellyfin:latest
    container_name: jellyfin
    runtime: nvidia
    environment:
      - ....
      - NVIDIA_VISIBLE_DEVICES=all
```

```
.....
```

4.) (Optional-Proxmox) Blacklist GPU auf pve-Host:

```
#AMD GPUs
echo "blacklist amdgpu" >> /etc/modprobe.d/blacklist.conf
echo "blacklist radeon" >> /etc/modprobe.d/blacklist.conf
#NVIDIA GPUs
echo "blacklist nouveau" >> /etc/modprobe.d/blacklist.conf
echo "blacklist nvidia*" >> /etc/modprobe.d/blacklist.conf
#Intel GPUs
echo "blacklist i915" >> /etc/modprobe.d/blacklist.conf
```

5.) Container verliert GPU ("Failed to initialize NVML: Unknown Error")

https://bobcares.com/blog/docker-failed-to-initialize-nvml-unknown-error/#:~:text=queries%20and%20issues.-,How%20to%20resolve%20docker%20failed%20to%20initialize%20NVML%20unknown%20error,and%20the%20GPUs%20return%20available.

https://github.com/NVIDIA/nvidia-docker/issues/1730

```
#Config editieren
nano /etc/nvidia-container-runtime/config.toml
# Uncomment folgende Zeile
no-cgroups = false
#Docker neu starten
sudo systemctl restart docker
#Testen mit test container
docker run -d --rm --runtime=nvidia --gpus all \
    --device=/dev/nvidia-uvm \
    --device=/dev/nvidia-uvm-tools \
    --device=/dev/nvidia-modeset \
    --device=/dev/nvidiactl \
    --device=/dev/nvidia0 \
    nvcr.io/nvidia/cuda:12.0.0-base-ubuntu20.04 bash -c "while [ true ]; do nvidia-smi -L;
sleep 5; done"
# Container ID
```

```
455dca9339e2184e3d0a93c1c216efa543642627783c1e0cbaf2c136162d89f9
# Log des Containers öffnen
docker logs 455dca9339e2184e3d0a93c1c216efa543642627783c1e0cbaf2c136162d89f9
# Absturz -Befehl testen
sudo systemctl daemon-reload
# Log wieder öffnen
docker logs 455dca9339e2184e3d0a93c1c216efa543642627783c1e0cbaf2c136162d89f9

#Falls kein Eintrag mit "Failed to initialize NVML: Unknown Error" auftaucht hat der Fix funktioniert
```

Revision #5
Created 2023-06-25 11:55:03 UTC by Dung
Updated 2023-07-29 13:49:00 UTC by Dung