Red Hat Virtualization 4.4

Setting up an NVIDIA GPU for a virtual machine in Red Hat Virtualization

How to configure a virtual machine in Red Hat Virtualization to use a dedicated GPU or vGPU.

Red Hat Virtualization 4.4 Setting up an NVIDIA GPU for a virtual machine in Red Hat Virtualization

How to configure a virtual machine in Red Hat Virtualization to use a dedicated GPU or vGPU.

Red Hat Virtualization Documentation Team
Red Hat Customer Content Services
rhev-docs@redhat.com

Legal Notice

Abstract

This document describes how to use a host with a graphics processing unit (GPU) to run virtual machines in Red Hat Virtualization for graphics-intensive tasks and software that cannot run without a GPU.

Table of Contents

PREFACE

You can use a host with a compatible graphics processing unit (GPU) to run virtual machines in Red Hat Virtualization that are suited for graphics-intensive tasks and for running software that cannot run without a GPU, such as CAD.

You can assign a GPU to a virtual machine in one of the following ways:

- **GPU passthrough**: You can assign a host GPU to a single virtual machine, so the virtual machine, instead of the host, uses the GPU.

- **Virtual GPU (vGPU)**: You can divide a physical GPU device into one or more virtual devices, referred to as *mediated devices*. You can then assign these mediated devices to one or more virtual machines as virtual GPUs. These virtual machines share the performance of a single physical GPU. For some GPUs, only one mediated device can be assigned to a single guest. vGPU support is only available on selected NVIDIA GPUs.

  ### Example:

  A host has four GPUs. Each GPU can support up to 16 vGPUs, for a total of 64 vGPUs. Some possible vGPU assignments are:

  - one virtual machine with 64 vGPUs

  - 64 virtual machines, each with one vGPU

  - 32 virtual machines, each with one vGPU; eight virtual machines, each with two vGPUs; 4 virtual machines, each with four vGPUs

CHAPTER 1. GPU DEVICE PASSTHROUGH: ASSIGNING A HOST GPU TO A SINGLE VIRTUAL MACHINE

Red Hat Virtualization supports PCI VFIO, also called device passthrough, for some NVIDIA PCIe-based GPU devices as non-VGA graphics devices.

You can attach one or more host GPUs to a single virtual machine by passing through the host GPU to the virtual machine, in addition to one of the standard emulated graphics interfaces. The virtual machine uses the emulated graphics device for pre-boot and installation, and the GPU takes control when its graphics drivers are loaded.

For information on the exact number of host GPUs that you can pass through to a single virtual machine, see the NVIDIA website.

To assign a GPU to a virtual machine, follow the steps in these procedures:

1. Enable the I/O Memory Management Unit (IOMMU) on the host machine.

2. Detach the GPU from the host.

3. Attach the GPU to the guest.

4. Install GPU drivers on the guest.

5. Configure Xorg on the guest.

These steps are detailed below.

Prerequisites

- Your GPU device supports GPU passthrough mode.

- Your system is listed as a validated server hardware platform.

- Your host chipset supports Intel VT-d or AMD-Vi

For more information about supported hardware and software, see Validated Platforms in the *NVIDIA GPU Software Release Notes*.

1.1. ENABLING HOST IOMMU SUPPORT AND BLACKLISTING NOUVEAU

I/O Memory Management Unit (IOMMU) support on the host machine is necessary to use a GPU on a virtual machine.

Procedure

1. In the Administration Portal, click **Compute → Hosts**. Select a host and click **Edit**. The **Edit Hosts** pane appears.

2. Click the **Kernel** tab.

3. Check the **Hostdev Passthrough & SR-IOV** checkbox. This checkbox enables IOMMU support for a host with Intel VT-d or AMD Vi by adding **intel_iommu=on** or **amd_iommu=on** to the kernel command line.

4. Check the **Blacklist Nouveau** checkbox.

5. Click **OK**.

6. Select the host and click **Management → Maintenance** and **OK**.

7. Click **Installation → Reinstall**.

8. After the reinstallation is finished, reboot the host machine.

9. When the host machine has rebooted, click **Management → Activate**.

> **NOTE**
>
> To enable IOMMU support using the command line, edit the **grub.conf** file in the virtual machine (./entries/rhvh-4.4.<machine id>.conf) to include the option **intel_iommu=on**.

1.2. DETACHING THE GPU FROM THE HOST

You cannot add the GPU to the virtual machine if the GPU is bound to the host kernel driver, so you must unbind the GPU device from the host before you can add it to the virtual machine. Host drivers often do not support dynamic unbinding of the GPU, so it is recommended to manually exclude the device from binding to the host drivers.

Procedure

1. On the host, identify the device slot name and IDs of the device by running the **lspci** command. In the following example, a graphics controller such as an NVIDIA Quadro or GRID card is used:

   ```
   # lspci -Dnn | grep -i NVIDIA
   0000:03:00.0 VGA compatible controller [0300]: NVIDIA Corporation GK104GL [Quadro K4200] [10de:11b4] (rev a1)
   0000:03:00.1 Audio device [0403]: NVIDIA Corporation GK104 HDMI Audio Controller [10de:0e0a] (rev a1)
   ```

   The output shows that the NVIDIA GK104 device is installed. It has a graphics controller and an audio controller with the following properties:

   - The device slot name of the graphics controller is **0000:03:00.0**, and the vendor-id:device-id for the graphics controller are **10de:11b4**.

   - The device slot name of the audio controller is **0000:03:00.1**, and the vendor-id:device-id for the audio controller are **10de:0e0a**.

2. Prevent the host machine driver from using the GPU device. You can use a vendor-id:device-id with the pci-stub driver. To do this, append the **pci-stub.ids** option, with the vendor-id:device-id as its value, to the **GRUB_CMDLINX_LINUX** environment variable located in the **/etc/sysconfig/grub** configuration file, for example:

   ```
   GRUB_CMDLINE_LINUX="crashkernel=auto resume=/dev/mapper/vg0-lv_swap rd.lvm.lv=vg0/lv_root rd.lvm.lv=vg0/lv_swap rhgb quiet intel_iommu=on pci-stub.ids=10de:11b4,10de:0e0a"
   ```

   When adding additional vendor IDs and device IDs for pci-stub, separate them with a comma.

3. Regenerate the boot loader configuration using grub2-mkconfig to include this option:

```
# grub2-mkconfig -o /etc/grub2.cfg
```

> **NOTE**
>
> When using a UEFI-based host, the target file should be **/etc/grub2-efi.cfg**.

4. Reboot the host machine.

5. Confirm that IOMMU is enabled, the host device is added to the list of pci-stub.ids, and Nouveau is blacklisted:

```
# cat /proc/cmdline
BOOT_IMAGE=(hd0,msdos1)/vmlinuz-4.18.0-147.el8.x86_64 root=/dev/mapper/vg0-lv_root
ro crashkernel=auto resume=/dev/mapper/vg0-lv_swap rd.lvm.lv=vg0/lv_root
rd.lvm.lv=vg0/lv_swap rhgb quiet intel_iommu=on ❶
pci-stub.ids=10de:11b4,10de:0e0a ❷
rdblacklist=nouveau ❸
```

❶ IOMMU is enabled

❷ the host device is added to the list of pci-stub.ids

❸ Nouveau is blacklisted

1.3. ATTACHING THE GPU TO A VIRTUAL MACHINE

After unbinding the GPU from host kernel driver, you can add it to the virtual machine and enable the correct driver.

Procedure

1. Follow the steps in Adding a Host Device to a Virtual Machine in the *Virtual Machine Management Guide*.

2. Run the virtual machine and log in to it.

3. Install the NVIDIA GPU driver on the virtual machine.

4. Verify that the correct kernel driver is in use for the GPU with the **lspci -nnk** command. For example:

```
# lspci -nnk
00:07.0 VGA compatible controller [0300]: NVIDIA Corporation GK104GL [Quadro K4200]
[10de:11b4] (rev a1)
Subsystem: Hewlett-Packard Company Device [103c:1096]
Kernel driver in use: nvidia
Kernel modules: nouveau, nvidia_drm, nvidia
```

1.4. INSTALLING THE GPU DRIVER ON THE VIRTUAL MACHINE

Procedure

1. Run the virtual machine and connect to it using the VNC or SPICE console.

2. Download the driver to the virtual machine. For information on getting the driver, see the Drivers page on the NVIDIA website.

3. Install the GPU driver.

> **IMPORTANT**
>
> **Linux only:** When installing the driver on a Linux guest operating system, you are prompted to update xorg.conf. If you do not update xorg.conf during the installation, you need to update it manually.

4. After the driver finishes installing, reboot the machine. **For Windows virtual machines,** fully power off the guest from the Administration portal or the VM portal, not from within the guest operating system.

> **IMPORTANT**
>
> **Windows only:** Powering off the virtual machine from within the Windows guest operating system sometimes sends the virtual machine into hibernate mode, which does not completely clear the memory, possibly leading to subsequent problems. Using the Administration portal or the VM portal to power off the virtual machine forces it to fully clean the memory.

5. Connect a monitor to the host GPU output interface and run the virtual machine.

6. Set up NVIDIA vGPU guest software licensing for each vGPU and add the license credentials in the NVIDIA control panel. For more information, see How NVIDIA vGPU Software Licensing Is Enforced in the *NVIDIA Virtual GPU Software Documentation*.

1.5. UPDATING AND ENABLING XORG (LINUX VIRTUAL MACHINES)

Before you can use the GPU on the virtual machine, you need to update and enable xorg on the virtual machine. The NVIDIA driver installation should do this automatically. Check if xorg is updated and enabled by viewing **/etc/X11/xorg.conf**:

```
# cat /etc/X11/xorg.conf
```

The first two lines say if it was generated by NVIDIA. For example:

```
# cat /etc/X11/xorg.conf
# nvidia-xconfig: X configuration file generated by nvidia-xconfig
# nvidia-xconfig: version 390.87 (buildmeister@swio-display-x64-rhel04-14) Tue Aug 21 17:33:38 PDT 2018
```

Procedure

1. On the virtual machine, generate the **xorg.conf** file using following command:

```
# X -configure
```

2. Copy the **xorg.conf** file to **/etc/X11/xorg.conf** using the following command:

```
# cp /root/xorg.conf.new /etc/X11/xorg.conf
```

3. Reboot the virtual machine.

4. Verify that xorg is updated and enabled by viewing **/etc/X11/xorg.conf**:

```
# cat /etc/X11/xorg.conf
```

Search for the **Device** section. You should see an entry similar to the following section:

```
Section "Device"
    Identifier    "Device0"
    Driver        "nvidia"
    VendorName    "NVIDIA Corporation"
EndSection
```

The GPU is now assigned to the virtual machine.

1.6. REMOVING A HOST GPU FROM A VIRTUAL MACHINE

- For information on removing a host GPU from a virtual machine, see Removing Host Devices from a Virtual Machine in the *Virtual Machine Management Guide*.

CHAPTER 2. ASSIGNING VIRTUAL GPUS

To set up NVIDIA vGPU devices, you need to:

1. Obtain and install the correct NVIDIA vGPU driver for your GPU device

2. Create mediated devices

3. Assign each mediated device to a virtual machine

4. Install guest drivers on each virtual machine.

The following procedures explain this process.

2.1. SETTING UP NVIDIA VGPU DEVICES ON THE HOST

> **NOTE**
>
> Before installing the NVIDIA vGPU driver on the guest operating system, you need to understand the licensing requirements and obtain the correct license credentials.

Prerequisites

- Your GPU device supports virtual GPU (vGPU) functionality.

- Your system is listed as a validated server hardware platform.

For more information about supported GPUs and validated platforms, see NVIDIA vGPU CERTIFIED SERVERS on www.nvidia.com.

Procedure

1. Download and install the NVIDIA–vGPU driver. For information on getting the driver, see vGPU drivers page on the NVIDIA website. An Nvidia enterprise account is required to download the drivers. Contact the hardware vendor if this is not available.

2. Unzip the downloaded file from the Nvidia website and copy it to the host to install the driver.

3. If the NVIDIA software installer did not create the **/etc/modprobe.d/nvidia-installer-disable-nouveau.conf** file, create it manually.

4. Open **/etc/modprobe.d/nvidia-installer-disable-nouveau.conf** file in a text editor and add the following lines to the end of the file:

   ```
   blacklist nouveau
   options nouveau modeset=0
   ```

5. Regenerate the initial ramdisk for the current kernel, then reboot:

   ```
   # dracut --force
   # reboot
   ```

   Alternatively, if you need to use a prior supported kernel version with mediated devices, regenerate the initial ramdisk for all installed kernel versions:

```
# dracut --regenerate-all --force
# reboot
```

6. Check that the kernel loaded the **nvidia_vgpu_vfio** module:

```
# lsmod | grep nvidia_vgpu_vfio
```

7. Check that the **nvidia-vgpu-mgr.service** service is running:

```
# systemctl status nvidia-vgpu-mgr.service
```

For example:

```
# lsmod | grep nvidia_vgpu_vfio
nvidia_vgpu_vfio 45011 0
nvidia 14333621 10 nvidia_vgpu_vfio
mdev 20414 2 vfio_mdev,nvidia_vgpu_vfio
vfio 32695 3 vfio_mdev,nvidia_vgpu_vfio,vfio_iommu_type1
# systemctl status nvidia-vgpu-mgr.service
nvidia-vgpu-mgr.service - NVIDIA vGPU Manager Daemon
   Loaded: loaded (/usr/lib/systemd/system/nvidia-vgpu-mgr.service; enabled; vendor preset:
disabled)
   Active: active (running) since Fri 2018-03-16 10:17:36 CET; 5h 8min ago
 Main PID: 1553 (nvidia-vgpu-mgr)
   [...]
```

8. In the Administration Portal, click **Compute → Virtual Machines**.

9. Click the name of the virtual machine to go to the details view.

10. Click the **Host Devices** tab.

11. Click **Manage vGPU**. The **Manage vGPU** dialog box opens.

12. Select a vGPU type and the number of instances that you would like to use with this virtual machine.

13. Select **On** for **Secondary display adapter for VNC** to add a second emulated QXL or VGA graphics adapter as the primary graphics adapter for the console in addition to the vGPU.

> **NOTE**
>
> On cluster levels 4.5 and later, when a vGPU is used and the **Secondary display adapter for VNC** is set to **On**, an additional framebuffer display device is automatically added to the virtual machine. This allows the virtual machine console to be displayed before the vGPU is initialized, instead of a blank screen.

14. Click **Save**.

## 2.2. INSTALLING THE VGPU DRIVER ON THE VIRTUAL MACHINE

Procedure

1. Run the virtual machine and connect to it using the VNC console.

   > **NOTE**
   >
   > SPICE is not supported on vGPU.

2. Download the driver to the virtual machine. For information on getting the driver, see the Drivers page on the NVIDIA website.

3. Install the vGPU driver, following the instructions in Installing the NVIDIA vGPU Software Graphics Driver in the *NVIDIA Virtual GPU software documentation*.

   > **IMPORTANT**
   >
   > **Linux only:** When installing the driver on a Linux guest operating system, you are prompted to update xorg.conf. If you do not update xorg.conf during the installation, you need to update it manually.

4. After the driver finishes installing, reboot the machine. **For Windows virtual machines,** fully power off the guest from the Administration portal or the VM portal, not from within the guest operating system.

   > **IMPORTANT**
   >
   > **Windows only:** Powering off the virtual machine from within the Windows guest operating system sometimes sends the virtual machine into hibernate mode, which does not completely clear the memory, possibly leading to subsequent problems. Using the Administration portal or the VM portal to power off the virtual machine forces it to fully clean the memory.

5. Run the virtual machine and connect to it using one of the supported remote desktop protocols, such as Mechdyne TGX, and verify that the vGPU is recognized by opening the NVIDIA Control Panel. On Windows, you can alternatively open the Windows Device Manager. The vGPU should appear under **Display adapters**. For more information, see the NVIDIA vGPU Software Graphics Driver in the *NVIDIA Virtual GPU software documentation*.

6. Set up NVIDIA vGPU guest software licensing for each vGPU and add the license credentials in the NVIDIA control panel. For more information, see How NVIDIA vGPU Software Licensing Is Enforced in the *NVIDIA Virtual GPU Software Documentation*.

2.3. REMOVING NVIDIA VGPU DEVICES

To change the configuration of assigned vGPU mediated devices, the existing devices have to be removed from the assigned guests.

Procedure

1. In the Administration Portal, click **Compute → Virtual Machines**.

2. Click the name of the virtual machine to go to the details view.

3. Click the **Host Devices** tab.

4. Click **Manage vGPU**. The **Manage vGPU** dialog box opens.

5. Click the **x** button next to **Selected vGPU Type Instances** to detach the vGPU from the virtual machine.

6. Click **SAVE**.

2.4. MONITORING NVIDIA VGPUS

For NVIDIA vGPUS, to get info on the physical GPU and vGPU, you can use the NVIDIA System Management Interface by entering the **nvidia-smi** command on the host. For more information, see NVIDIA System Management Interface nvidia-smi in the *NVIDIA Virtual GPU Software Documentation*.

For example:

```
# nvidia-smi
Thu Nov  1 17:40:09 2018
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 410.62             Driver Version: 410.62                 |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla M60         On   | 00000000:84:00.0 Off |           Off |
| N/A  40C   P8   24W / 150W |   1034MiB /  8191MiB |     0%     Default |
+-------------------------------+----------------------+----------------------+
|   1  Tesla M60         On   | 00000000:85:00.0 Off |           Off |
| N/A  33C   P8   23W / 150W |   8146MiB /  8191MiB |     0%     Default |
+-------------------------------+----------------------+----------------------+
|   2  Tesla M60         On   | 00000000:8B:00.0 Off |           Off |
| N/A  34C   P8   24W / 150W |   8146MiB /  8191MiB |     0%     Default |
+-------------------------------+----------------------+----------------------+
|   3  Tesla M60         On   | 00000000:8C:00.0 Off |           Off |
| N/A  45C   P8   24W / 150W |     18MiB /  8191MiB |     0%     Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                         GPU Memory |
|  GPU       PID   Type   Process name                   Usage      |
|=============================================================================|
|   0    34432   C+G   vgpu                             508MiB |
|   0    34718   C+G   vgpu                             508MiB |
|   1    35032   C+G   vgpu                            8128MiB |
|   2    35032   C+G   vgpu                            8128MiB |
+-----------------------------------------------------------------------------+
```

2.5. REMOTE DESKTOP STREAMING SERVICES FOR NVIDIA VGPU

The following remote desktop streaming services have been successfully tested for use with the NVIDIA vGPU feature in RHEL 8:

- HP-RGS

- Mechdyne TGX – It is currently not possible to use Mechdyne TGX with Windows Server 2016 guests.

- NICE DCV – When using this streaming service, use fixed resolution settings, because using dynamic resolution in some cases results in a black screen.

CHAPTER 3. RELATED INFORMATION

For further information on using NVIDA vGPU on RHEL with KVM, see:

- the NVIDIA GPU Software Release Notes .

- *NVIDIA Virtual GPU Software Documentation* at https://docs.nvidia.com.

APPENDIX A. LEGAL NOTICE