

NVIDIA Virtual GPU for KVM Guests

Publication Date: September 28, 2024

Contents

- 1 Introduction 2
- 2 Configuring vGPU manager in VM Host Server 2
- 3 Creating a vGPU device 5
- 4 Assign the vGPU device to a VM Guest 13
- 5 Configuring vGPU in VM Guest 14
- 6 Licensing vGPU in the VM Guest 15
- 7 Configuring a graphics mode 17
- 8 Configuring compute mode 18
- 9 Additional tasks 19
- 10 For more information 21
- 11 NVIDIA virtual GPU background 22
- 12 GNU Free Documentation License 25

1 Introduction

NVIDIA virtual GPU (vGPU) is a graphics virtualization solution that provides multiple virtual machines (VMs) simultaneous access to one physical Graphics Processing Unit (GPU) on the VM Host Server. This article refers to the Volta and Ampere GPU architecture.

2 Configuring vGPU manager in VM Host Server

2.1 Prepare VM Host Server environment

1. Verify that you have a compatible server and GPU cards. Check specifications for details:

- <https://docs.nvidia.com/grid/gpus-supported-by-vgpu.html> ↗
- <https://docs.nvidia.com/grid/index.html> ↗

2. Verify that VM Host Server is SUSE Linux Enterprise Server 15 SP3 or newer:

```
> cat /etc/issue
Welcome to SUSE Linux Enterprise Server 15 SP3 (x86_64) - Kernel \r (\l).
```

3. Get the vGPU drivers from NVIDIA. In order to get the software, please follow the steps at <https://docs.nvidia.com/grid/latest/grid-software-quick-start-guide/index.html#redeeming-pak-and-downloading-grid-software> ↗. For example, for vGPU 13.0 installation, you will need the following files:

```
NVIDIA-Linux-x86_64-470.63-vgpu-kvm.run # vGPU manager for the VM host
NVIDIA-Linux-x86_64-470.63.01-grid.run # vGPU driver for the VM guest
```

4. If you are using Ampere architecture GPU cards, verify that VM Host Server supports VT-D/IOMMU and SR-IOV technologies, and that they are enabled in BIOS.
5. Enable IOMMU. Verify that it is included in the boot command line:

```
cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-default [...] intel_iommu=on [...]
```

If not, add the following line to [/etc/default/grub](#).

- For Intel CPUs:

```
GRUB_CMDLINE_LINUX="intel_iommu=on"
```

For AMD CPUs:

```
GRUB_CMDLINE_LINUX="amd_iommu=on"
```

Then generate new GRUB 2 configuration file and reboot:

```
> sudo grub2-mkconfig -o /boot/grub2/grub.cfg  
> sudo systemctl reboot
```



Tip

You can verify that IOMMU is loaded by running the following command:

```
sudo dmesg | grep -e IOMMU
```

6. Enable SR-IOV. Refer to <https://docs.nvidia.com/grid/13.0/grid-vgpu-user-guide/index.html#vgpu-types-tesla-v100-pcie> for useful information.
7. Disable the nouveau kernel module by adding the following line it to the top of the `/etc/modprobe.d/50-blacklist.conf` file:

```
blacklist nouveau
```

2.2 Install the NVIDIA KVM driver

1. Exit from the graphical mode:

```
> sudo init 3
```

2. Install `kernel-default-devel` and `gcc` packages and their dependencies:

```
> sudo zypper in kernel-default-devel gcc
```

3. Download the vGPU software from the NVIDIA portal. Make the NVIDIA vGPU driver executable and run it:

```
> chmod +x NVIDIA-Linux-x86_64-450.55-vgpu-kvm.run
> sudo ./NVIDIA-Linux-x86_64-450.55-vgpu-kvm.run
```

You can find detailed information about the installation process in the log file `/var/log/nvidia-installer.log`



Tip

To enable dynamic kernel-module support, and thus have the module rebuilt automatically when new a new kernel is installed, add the `--dkms` option:

```
> sudo ./NVIDIA-Linux-x86_64-450.55-vgpu-kvm.run --dkms
```

4. When the driver installation is finished, reboot the system:

```
> sudo systemctl reboot
```

2.3 Verify the driver installation

1. Verify loaded kernel modules:

```
> lsmod | grep nvidia
nvidia_vgpu_vfio      49152  9
nvidia                14393344  229 nvidia_vgpu_vfio
mdev                  20480  2 vfio_mdev,nvidia_vgpu_vfio
vfio                   32768  6 vfio_mdev,nvidia_vgpu_vfio,vfio_iommu_type1
```

The modules containing the `vfio` string are required dependencies.

2. Print the GPU device status with the `nvidia-smi` command. The output should be similar to the following one:

```
> nvidia-smi
+-----+
| NVIDIA-SMI 470.63      Driver Version: 470.63      CUDA Version: N/A      |
+-----+-----+-----+-----+
| 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	A40	Off	00000000:31:00.0	Off	0
0%	46C	P0	39W / 300W	0MiB / 45634MiB	0%	Default
						N/A

+-----						
Processes:						
GPU	GI	CI	PID	Type	Process name	GPU Memory Usage
ID	ID	ID				
=====						
No running processes found						

+-----						

3. Check the sysfs file system. For Volta and earlier GPU cards, new directory `mdev_supported_types` is added, for example:

```
cd /sys/bus/pci/devices/00000000\:31\:00.0/mdev_supported_types
```

For Ampere GPU cards, the directory will be created automatically for each virtual function after SR-IOV is enabled.

3 Creating a vGPU device

3.1 Create a legacy vGPU device without support for SR-IOV

All the NVIDIA Volta and earlier architecture GPUs work in this mode.

1. Obtain the Bus/Device/Function (BDF) numbers of the host GPU device:

```
> lspci | grep NVIDIA
84:00.0 3D controller: NVIDIA Corporation GV100GL [Tesla V100 PCIe 16GB] (rev a1)
```

2. Check for the mdev supported devices and detailed information:

```
> ls /sys/bus/pci/devices/0000:84:00.0/mdev_supported_types/
nvidia-105 nvidia-106 nvidia-107 nvidia-108 nvidia-109 nvidia-110 [...]
```

The map of vGPU mdev devices and their type is as follows:

- nvidia-105 to nvidia-109: 1Q 2Q 4Q 8Q 16Q
- nvidia-110 to nvidia-114: 1A 2A 4A 8A 16A
- nvidia-115, nvidia-163, nvidia-217, nvidia-247: 1B 2B 2B4 1B4
- nvidia-299 to nvidia-301: 4C 8C 16C

Refer to <https://docs.nvidia.com/grid/latest/grid-vgpu-user-guide/index.html#vgpu-types-tesla-v100-pcie> for more details.

3. Inspect a vGPU device:

```
> cd /sys/bus/pci/devices/0000:03:00.0/mdev_supported_types/  
> ls nvidia-105  
> cat nvidia-105/description  
num_heads=2, frl_config=60, framebuffer=1024M, max_resolution=4096x2160,  
max_instance=16  
> cat nvidia-105/name  
GRID V100-1Q
```

4. Generate a unique ID and create an mdev device based on it:

```
> uuidgen  
4f3b6e47-0baa-4900-b0b1-284c1ecc192f  
> sudo echo "4f3b6e47-0baa-4900-b0b1-284c1ecc192f" > nvidia-105/create
```

5. Verify the new mdev device. You can inspect the content of the /sys/bus/mdev/devices directory:

```
> cd /sys/bus/mdev/devices  
> ls -l  
lrwxrwxrwx 1 root root 0 Aug 30 23:03 86380ffb-8f13-4685-9c48-0e0f4e65fb87 \  
-> ../../../../devices/  
pci0000:80/0000:80:02.0/0000:84:00.0/86380ffb-8f13-4685-9c48-0e0f4e65fb87  
lrwxrwxrwx 1 root root 0 Aug 30 23:03 86380ffb-8f13-4685-9c48-0e0f4e65fb88 \  
-> ../../../../devices/  
pci0000:80/0000:80:02.0/0000:84:00.0/86380ffb-8f13-4685-9c48-0e0f4e65fb88  
lrwxrwxrwx 1 root root 0 Aug 30 23:03 86380ffb-8f13-4685-9c48-0e0f4e65fb89 \  
-> ../../../../devices/  
pci0000:80/0000:80:02.0/0000:84:00.0/86380ffb-8f13-4685-9c48-0e0f4e65fb89  
lrwxrwxrwx 1 root root 0 Aug 30 23:03 86380ffb-8f13-4685-9c48-0e0f4e65fb90 \  
-> ../../../../devices/
```

```
-> ../../../../devices/  
pci0000:80/0000:80:02.0/0000:84:00.0/86380ffb-8f13-4685-9c48-0e0f4e65fb90
```

Or you can use the `mdevctl` command:

```
> sudo mdevctl list  
86380ffb-8f13-4685-9c48-0e0f4e65fb90 0000:84:00.0 nvidia-299  
86380ffb-8f13-4685-9c48-0e0f4e65fb89 0000:84:00.0 nvidia-299  
86380ffb-8f13-4685-9c48-0e0f4e65fb87 0000:84:00.0 nvidia-299  
86380ffb-8f13-4685-9c48-0e0f4e65fb88 0000:84:00.0 nvidia-299
```

6. Query the new vGPU device capability:

```
> sudo nvidia-smi vgpu -q  
GPU 00000000:84:00.0  
Active vGPUs : 1  
vGPU ID : 3251634323  
  VM UUID : ee7b7a4b-388a-4357-a425-5318b2c65b3f  
  VM Name : sle15sp3  
  vGPU Name : GRID V100-4C  
  vGPU Type : 299  
  vGPU UUID : d471c7f2-0a53-11ec-afd3-38b06df18e37  
  MDEV UUID : 86380ffb-8f13-4685-9c48-0e0f4e65fb87  
  Guest Driver Version : 460.91.03  
  License Status : Licensed  
  GPU Instance ID : N/A  
  Accounting Mode : Disabled  
  ECC Mode : N/A  
  Accounting Buffer Size : 4000  
  Frame Rate Limit : N/A  
  FB Memory Usage  
    Total : 4096 MiB  
    Used : 161 MiB  
    Free : 3935 MiB  
  Utilization  
    Gpu : 0 %  
    Memory : 0 %  
    Encoder : 0 %  
    Decoder : 0 %  
  Encoder Stats  
    Active Sessions : 0  
    Average FPS : 0  
    Average Latency : 0  
  FBC Stats  
    Active Sessions : 0  
    Average FPS : 0  
    Average Latency : 0
```

3.2 Create a vGPU device with support for SR-IOV

All NVIDIA Ampere and newer architecture GPUs work in this mode.

1. Obtain the Bus/Device/Function (BDF) numbers of the host GPU device:

```
> lspci | grep NVIDIA
b1:00.0 3D controller: NVIDIA Corporation GA100 [A100 PCIe 40GB] (rev a1)
```

2. Enable virtual functions:

```
> sudo /usr/lib/nvidia/sriov-manage -e 00:b1:0000.0
```



Note

This configuration is not persistent and must be re-enabled after the host reboot.

3. Obtain the Bus/Domain/Function (BDF) of virtual functions on the GPU:

```
> ls -l /sys/bus/pci/devices/0000:b1:00.0/ | grep virtfn
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn0 -> ../0000:b1:00.4
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn1 -> ../0000:b1:00.5
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn10 -> ../0000:b1:01.6
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn11 -> ../0000:b1:01.7
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn12 -> ../0000:b1:02.0
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn13 -> ../0000:b1:02.1
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn14 -> ../0000:b1:02.2
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn15 -> ../0000:b1:02.3
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn2 -> ../0000:b1:00.6
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn3 -> ../0000:b1:00.7
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn4 -> ../0000:b1:01.0
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn5 -> ../0000:b1:01.1
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn6 -> ../0000:b1:01.2
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn7 -> ../0000:b1:01.3
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn8 -> ../0000:b1:01.4
lrwxrwxrwx 1 root root      0 Sep 21 11:58 virtfn9 -> ../0000:b1:01.5
```

4. *Create a vGPU device.* Select the virtual function (VF) that you want to use to create the vGPU device and assign it a unique ID.



Important

Each VF can only create one vGPU instance. If you want to create more vGPU instances, you need to use a different VF.

```
> cd /sys/bus/pci/devices/0000:b1:00.0/virtfn1/mdev_supported_types
> for i in *; do echo "$i" $(cat $i/name) available: $(cat $i/avail*); done
nvidia-468 GRID A100-4C available: 0
nvidia-469 GRID A100-5C available: 0
nvidia-470 GRID A100-8C available: 0
nvidia-471 GRID A100-10C available: 1
nvidia-472 GRID A100-20C available: 0
nvidia-473 GRID A100-40C available: 0
nvidia-474 GRID A100-1-5C available: 0
nvidia-475 GRID A100-2-10C available: 0
nvidia-476 GRID A100-3-20C available: 0
nvidia-477 GRID A100-4-20C available: 0
nvidia-478 GRID A100-7-40C available: 0
nvidia-479 GRID A100-1-5CME available: 0
> uuidgen
f715f63c-0d00-4007-9c5a-b07b0c6c05de
> sudo echo "f715f63c-0d00-4007-9c5a-b07b0c6c05de" > nvidia-471/create
> sudo dmesg | tail
[...]
[ 3218.491843] vfio_mdev f715f63c-0d00-4007-9c5a-b07b0c6c05de: Adding to iommu group
322
[ 3218.499700] vfio_mdev f715f63c-0d00-4007-9c5a-b07b0c6c05de: MDEV: group_id = 322
[ 3599.608540] vfio_mdev f715f63c-0d00-4007-9c5a-b07b0c6c05de: Removing from iommu
group 322
[ 3599.616753] vfio_mdev f715f63c-0d00-4007-9c5a-b07b0c6c05de: MDEV: detaching iommu
[ 3626.345530] vfio_mdev f715f63c-0d00-4007-9c5a-b07b0c6c05de: Adding to iommu group
322
[ 3626.353383] vfio_mdev f715f63c-0d00-4007-9c5a-b07b0c6c05de: MDEV: group_id = 322
```

5. Verify the new vGPU device:

```
> cd /sys/bus/mdev/devices/
> ls
f715f63c-0d00-4007-9c5a-b07b0c6c05de
```

6. Query the new vGPU device capability:

```
> sudo nvidia-smi vgpu -q
GPU 00000000:B1:00.0
```

```

Active vGPUs           : 1
vGPU ID               : 3251634265
  VM UUID             : b0d9f0c6-a6c2-463e-967b-06cb206415b6
  VM Name             : sles15sp2-gehc-vm1
  vGPU Name           : GRID A100-10C
  vGPU Type           : 471
  vGPU UUID           : 444f610c-1b08-11ec-9554-ebd10788ee14
  MDEV UUID           : f715f63c-0d00-4007-9c5a-b07b0c6c05de
  Guest Driver Version : N/A
  License Status       : N/A
  GPU Instance ID     : N/A
  Accounting Mode      : N/A
  ECC Mode             : Disabled
  Accounting Buffer Size : 4000
  Frame Rate Limit     : N/A
  FB Memory Usage
    Total              : 10240 MiB
    Used               : 0 MiB
    Free               : 10240 MiB
  Utilization
    Gpu                : 0 %
    Memory             : 0 %
    Encoder             : 0 %
    Decoder             : 0 %
  Encoder Stats
    Active Sessions    : 0
    Average FPS        : 0
    Average Latency    : 0
  FBC Stats
    Active Sessions    : 0
    Average FPS        : 0
    Average Latency    : 0

```

3.3 Creating a MIG-backed vGPU



Important

SR-IOV is required to be enabled if you want to create vGPUs and assign them to guest VMs.

1. Enable MIG mode for a GPU:

```
> sudo nvidia-smi -i 0 -mig 1
```

```
Enabled MIG Mode for GPU 00000000:B1:00.0
All done.
```

2. Query the GPU instance profile:

```
> sudo nvidia-smi mig -lgip
+-----+
| GPU instance profiles:                                     |
| GPU   Name                ID   Instances   Memory   P2P   SM   DEC   ENC   |
|                               Free/Total   GiB           CE   JPEG  OFA   |
|=====|
|   0   MIG 1g.5gb          19    7/7        4.75    No    14   0    0   |
|                               |           |           |     |     |     |
|                               |           |           |     |     |     |
|                               |           |           |     |     |     |
+-----+
|   0   MIG 1g.5gb+me       20    1/1        4.75    No    14   1    0   |
|                               |           |           |     |     |     |
|                               |           |           |     |     |     |
|                               |           |           |     |     |     |
+-----+
|   0   MIG 2g.10gb         14    3/3        9.75    No    28   1    0   |
|                               |           |           |     |     |     |
|                               |           |           |     |     |     |
|                               |           |           |     |     |     |
+-----+
|   0   MIG 3g.20gb         9     2/2       19.62   No    42   2    0   |
|                               |           |           |     |     |     |
|                               |           |           |     |     |     |
|                               |           |           |     |     |     |
+-----+
|   0   MIG 4g.20gb         5     1/1       19.62   No    56   2    0   |
|                               |           |           |     |     |     |
|                               |           |           |     |     |     |
|                               |           |           |     |     |     |
+-----+
|   0   MIG 7g.40gb         0     1/1       39.50   No    98   5    0   |
|                               |           |           |     |     |     |
|                               |           |           |     |     |     |
|                               |           |           |     |     |     |
+-----+
```

3. Create a GPU instance specifying '5' as a GPU profile instance ID and optionally create a Compute Instance on it, either on the host server or within the guest:

```
> sudo nvidia-smi mig -cgi 5
Successfully created GPU instance ID 1 on GPU 0 using profile MIG 4g.20gb (ID 5)
> sudo nvidia-smi mig -cci -gi 1
Successfully created compute instance ID 0 on GPU 0 GPU instance ID 1 using
profile MIG 4g.20gb (ID 3)
```

4. Verify the GPU instance:

```
> sudo nvidia-smi
Tue Sep 21 11:19:36 2021
+-----+
| NVIDIA-SMI 470.63          Driver Version: 470.63          CUDA Version: N/A          |
|-----+-----+-----+-----+-----+-----+-----+
| GPU   Name                Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
+-----+-----+-----+-----+-----+-----+-----+
| 0     MIG 4g.20gb         0           | 00000000:01:00.0  0     | N                   |
+-----+-----+-----+-----+-----+-----+-----+
| 1     MIG 4g.20gb         0           | 00000000:01:00.0  0     | N                   |
+-----+-----+-----+-----+-----+-----+-----+
| 2     MIG 4g.20gb         0           | 00000000:01:00.0  0     | N                   |
+-----+-----+-----+-----+-----+-----+-----+
```

```

| Fan  Temp  Perf  Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
|                               |                     |             MIG M. |
|====+=====+=====+=====|
|  0  NVIDIA A100-PCI... On | 00000000:B1:00.0 Off |             On |
| N/A  38C   P0   38W / 250W |    0MiB / 40536MiB |           N/A   Default |
|                               |                     |             Enabled |
+-----+-----+-----+-----+
|
| MIG devices:
|
+-----+-----+-----+-----+
| GPU  GI  CI  MIG |      Memory-Usage |      Vol |      Shared |
|      ID ID  Dev |      BAR1-Usage | SM   Unc| CE  ENC  DEC  OFA  JPG|
|                               |                     |      ECC|          |
|====+=====+=====+=====|
|  0   1   0   0 |    0MiB / 20096MiB | 56    0 | 4   0   2   0   0 |
|                               |    0MiB / 32767MiB |         |          |
+-----+-----+-----+-----+
|
| Processes:
| GPU  GI  CI      PID  Type  Process name                       GPU Memory |
|      ID ID                                     Usage |
|====+=====+=====+=====|
| No running processes found |
+-----+-----+-----+-----+

```

5. Use the MIG instance. You can use the instance directly with the UUID—for example, assign it to a container or CUDA process.

You can also create a vGPU on top of it and assign it to a VM guest. The procedure is the same as for the vGPU with SR-IOV support. Refer to *Section 3.2, “Create a vGPU device with support for SR-IOV”*.

```

> sudo nvidia-smi -L
GPU 0: NVIDIA A100-PCI-E-40GB (UUID: GPU-ee14e29d-dd5b-2e8e-eeaf-9d3debd10788)
MIG 4g.20gb Device 0: (UUID: MIG-fed03f85-fd95-581b-837f-d582496d0260)

```

4 Assign the vGPU device to a VM Guest

4.1 Assign by libvirt

1. Create a `libvirt`-based virtual machine (VM) with UEFI support and a normal VGA display.
2. Edit the VM's configuration by running `virsh edit VM-NAME`.
3. Add the new mdev device with the unique ID you used when creating the vGPU device to the `<devices/>` section.



Note

If you are using Q-series, use `display='on'` instead.

```
<hostdev mode='subsystem' type='mdev' managed='no' model='vfio-pci' display='off'>
  <source>
    <address uuid='4f3b6e47-0baa-4900-b0b1-284c1ecc192f' />
  </source>
  <address type='pci' domain='0x0000' bus='0x00' slot='0x0a' function='0x0' />
</hostdev>
```

4.2 Assign by QEMU

Add the following device to the QEMU command line. Use the unique ID that you used when creating the vGPU device:

```
-device vfio-pci,sysfsdev=/sys/bus/mdev/devices/4f3b6e47-0baa-4900-b0b1-284c1ecc192f
```

5 Configuring vGPU in VM Guest

5.1 Prepare the VM Guest

- During VM Guest installation, disable secure boot, enable the SSH service, and select `wicked` for networking.
- Disable the `nouveau` video driver. Edit the file `/etc/modprobe.d/50-blacklist.conf` and add the following line to its upper section:

```
blacklist nouveau
```



Important

Disabling `nouveau` will work after you re-generate the `initrd` image with `mkinitrd`, and then reboot the VM Guest.

5.2 Install the vGPU driver in the VM Guest

1. Install the following packages and their dependencies:

```
> sudo zypper install kernel-default-devel libglvnd-devel
```

2. Download the vGPU software from the NVIDIA portal. Make the NVIDIA vGPU driver executable and run it:

```
> chmod +x NVIDIA-Linux-x86_64-470.63.01-grid.run  
> sudo ./NVIDIA-Linux-x86_64-470.63.01-grid.run
```



Tip

To enable dynamic kernel module support in order to get the module rebuilt automatically when new a new kernel is installed, add the `--dkms` option:

```
> sudo ./NVIDIA-Linux-x86_64-470.63.01-grid.run --dkms
```

3. During driver installation, select to run the `nvidia-xconfig` utility.

4. Verify the driver installation by checking the output of the `nvidia-smi` command:

```
> sudo nvidia-smi
+-----+
| NVIDIA-SMI 470.63.01    Driver Version: 470.63.01    CUDA Version: 11.4    |
+-----+-----+
| 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   GRID A100-10C         On   | 00000000:07:00.0 Off |             0      |
| N/A   N/A    P0     N/A /  N/A |  930MiB / 10235MiB |      0%    Default  |
|                                       |                  |     Disabled |
+-----+-----+-----+

+-----+
| Processes:                                                       |
| GPU  GI    CI          PID   Type   Process name                      GPU Memory |
|      ID    ID                                   |            | Usage   |
+-----+-----+-----+
| No running processes found                                     |
+-----+
```

6 Licensing vGPU in the VM Guest

1. Create the configuration file `/etc/nvidia/gridd.conf` based on `/etc/nvidia/gridd.conf.template`.
2. a. For licenses that are served from the NVIDIA License System, update the following options:

FeatureType

For GPU passthrough, set `FeatureType` to `4` for computing and `2` for graphic purposes. In case of a virtual GPU, whatever vGPU type is created via `mdev` determines the feature set that is enabled in VM Guest.

ClientConfigTokenPath

Optional: If you want to store the client configuration token in a custom location, add the `ClientConfigTokenPath` configuration parameter on a new line as `ClientConfigTokenPath="PATH_TO_TOKEN"`. By default, the client searches for the client configuration token in the `/etc/nvidia/ClientConfigToken/` directory.

Copy the client configuration token to the directory in which you want to store it.

- b. For licenses that are served from the legacy NVIDIA vGPU software license server, update the following options:

ServerAddress

Add your license server IP address.

ServerPort

Use the default "7070" or the port configured during the server setup.

FeatureType

For GPU passthrough, set `FeatureType` to `4` for computing and `2` for graphic purposes. In case of a virtual GPU, whatever vGPU type is created via `mdev` determines the feature set that is enabled in VM Guest.

3. Restart the `nvidia-gridd` service:

```
> sudo systemctl restart nvidia-gridd.service
```

4. Inspect the log file for possible errors:

```
> sudo grep gridd /var/log/messages
[...]
Aug 5 15:40:06 localhost nvidia-gridd: Started (4293)
Aug 5 15:40:24 localhost nvidia-gridd: License acquired successfully.
```


7 Configuring a graphics mode

7.1 Create or update the `/etc/X11/xorg.conf` file

1. If there is no `/etc/X11/xorg.conf` on the VM Guest, run the `nvidia-xconfig` utility.
2. Query the GPU device for detailed information:

```
> nvidia-xconfig --query-gpu-info
Number of GPUs: 1

GPU #0:
Name      : GRID V100-16Q
UUID      : GPU-089f39ad-01cb-11ec-89dc-da10f5778138
PCI BusID : PCI:0:10:0

Number of Display Devices: 0
```

3. Add GPU's BusID to `/etc/X11/xorg.conf`, for example:

```
Section "Device"
Identifier "Device0"
Driver "nvidia"
BusID "PCI:0:10:0"
VendorName "NVIDIA Corporation"
EndSection
```

7.2 Verify the graphics mode

Verify the following:

- A graphic desktop is booted correctly.
- The 'X' process of a running X-server is running in GPU:

```
> nvidia-smi
+-----+
| NVIDIA-SMI 470.63.01    Driver Version: 470.63.01    CUDA Version: 11.4    |
+-----+-----+-----+-----+
| 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. |
+-----+-----+-----+-----+
```

GPU	GI	CI	PID	Type	Process name	GPU Memory Usage
0	N/A	N/A	1921	G	/usr/bin/X	76MiB
0	N/A	N/A	1957	G	/usr/bin/gnome-shell	87MiB

7.3 Remote display

You need to install and configure the VNC server package `x11vnc` inside the VM Guest, and start it with the following command:

```
> sudo x11vnc -display :0 -auth /run/user/1000/gdm/Xauthority -forever -shared -ncache -bg -usepw -geometry 1900x1080
```

You can use `virt-manager` or `virt-viewer` to display the graphical output of a VM Guest.

! Important

For a `libvirt`-based VM Guest, verify that its XML configuration includes `display=on` as suggested in *Section 4.1, "Assign by libvirt"*.

8 Configuring compute mode

1. Download and install the CUDA toolkit. You can find it at https://developer.nvidia.com/cuda-downloads?target_os=Linux&target_arch=x86_64&Distribution=SLES&target_version=15&target_type=runfile_local.
2. Download CUDA samples from <https://github.com/nvidia/cuda-samples>.
3. Run CUDA sampling example:

```
> cd YOUR_GIT_CLONE_LOCATION/cuda-samples/Samples/0_Introduction/clock
```

```

> make
/usr/local/cuda/bin/nvcc -ccbin g++ -I../common/inc -m64 --threads 0 -gencode
arch=compute_35,code=sm_35 -gencode arch=compute_37,code=sm_37 -gencode
[...]
mkdir -p ../../bin/x86_64/linux/release
cp clock ../../bin/x86_64/linux/release
> ./clock
CUDA Clock sample
GPU Device 0: "Volta" with compute capability 7.0
Average clocks/block = 2820.718750

```

9 Additional tasks

This section introduces additional procedures that may be helpful after you have configured your vGPU.

9.1 Disabling Frame Rate Limiter

Frame Rate Limiter (FRL) is enabled by default. It limits the vGPU to a fixed frame rate, for example 60fps. If you experience a bad graphic display, you may need to disable FRL, for example:

```

> sudo echo "frame_rate_limiter=0" > /sys/bus/mdev/
devices/86380ffb-8f13-4685-9c48-0e0f4e65fb87/nvidia/vgpu_params

```

9.2 Enabling/Disabling Error Correcting Code (ECC)

Since the NVIDIA Pascal architecture, NVIDIA GPU Cards support ECC memory to improve data integrity. ECC is also supported by software since NVIDIA vGPU 9.0.

To enable ECC:

```

> sudo nvidia-smi -e 1
> nvidia-smi -q
Ecc Mode
  Current           : Enabled
  Pending           : Enabled

```

To disable ECC:

```

> sudo nvidia-smi -e 0

```

9.3 Black screen in Virt-manager

If you see only a black screen in Virt-manager, press **Alt - Ctrl - 2** from Virt-manager viewer. You should be able to get in the display again.

9.4 Black screen in VNC client when using a non-QEMU VNC server

Use the `xvnc` server.

9.5 Kernel panic occurs because the Nouveau and NVIDIA drivers compete on GPU resources

The boot messages will look as follows:

```
[ 16.742439] Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.14.0-0-
g155821a1990b-prebuilt.qemu.org 04/01/2014
[ 16.742441] RIP: 0010:__pci_enable_msi_range+0x3a9/0x3f0
[ 16.742443] Code: 76 60 49 8d 56 50 48 89 df e8 73 f6 fc ff e9 3b fe ff ff 31 f6 48 89
df e8 64 73 fd ff e9 d6 fe ff ff 44 89 fd e9 1a ff ff ff <0f> 0b bd ea ff ff ff e9 0e ff
ff ff bd ea ff ff ff e9 04 ff f
f ff
[ 16.742444] RSP: 0018:ffffb04bc052fb28 EFLAGS: 00010202
[ 16.742445] RAX: 0000000000000010 RBX: ffff9e93a85bc000 RCX: 0000000000000001
[ 16.742457] RDX: 0000000000000000 RSI: 0000000000000001 RDI: ffff9e93a85bc000
[ 16.742458] RBP: ffff9e93a2550800 R08: 0000000000000002 R09: fffffb04bc052fb1c
[ 16.742459] R10: 0000000000000050 R11: 0000000000000020 R12: ffff9e93a2550800
[ 16.742459] R13: 0000000000000001 R14: ffff9e93a2550ac8 R15: 0000000000000001
[ 16.742460] FS: 00007f9f26889740(0000) GS:ffff9e93bfdc0000(0000) knlGS:0000000000000000
[ 16.742461] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033
[ 16.742462] CR2: 000000000008aeb90 CR3: 0000000286470003 CR4: 000000000170ee0
[ 16.742465] Call Trace:
[ 16.742503] ? __pci_find_next_cap_ttl+0x93/0xd0
[ 16.742505] pci_enable_msi+0x16/0x30
[ 16.743039] nv_init_msi+0x1a/0xf0 [nvidia]
[ 16.743154] nv_open_device+0x81b/0x890 [nvidia]
[ 16.743248] nvidia_open+0x2f7/0x4d0 [nvidia]
[ 16.743256] ? kobj_lookup+0x113/0x160
[ 16.743354] nvidia_frontend_open+0x53/0x90 [nvidia]
[ 16.743361] chrdev_open+0xc4/0x1a0
[ 16.743370] ? cdev_put.part.2+0x20/0x20
[ 16.743374] do_dentry_open+0x204/0x3a0
[ 16.743378] path_openat+0x2fc/0x1520
[ 16.743382] ? unlazy_walk+0x32/0xa0
```

```

[ 16.743383] ? terminate_walk+0x8c/0x100
[ 16.743385] do_filp_open+0x9b/0x110
[ 16.743387] ? chown_common+0xf7/0x1c0
[ 16.743390] ? kmem_cache_alloc+0x18a/0x270
[ 16.743392] ? do_sys_open+0x1bd/0x260
[ 16.743394] do_sys_open+0x1bd/0x260
[ 16.743400] do_syscall_64+0x5b/0x1e0
[ 16.743409] entry_SYSCALL_64_after_hwframe+0x44/0xa9
[ 16.743418] RIP: 0033:0x7f9f2593961d
[ 16.743420] Code: f0 25 00 00 41 00 3d 00 00 41 00 74 48 64 8b 04 25 18 00 00 00 85 c0
75 64 89 f2 b8 01 01 00 00 48 89 fe bf 9c ff ff ff 0f 05 <48> 3d 00 f0 ff ff 0f 87 97 00
00 00 48 8b 4c 24 28 64 48 33 0
c 25
[ 16.743420] RSP: 002b:00007ffcfa214930 EFLAGS: 00000246 ORIG_RAX: 0000000000000101
[ 16.743422] RAX: ffffffffda RBX: 00007ffcfa214c30 RCX: 00007f9f2593961d
[ 16.743422] RDX: 00000000000008002 RSI: 00007ffcfa2149b0 RDI: 00000000fffff9c
[ 16.743423] RBP: 00007ffcfa2149b0 R08: 0000000000000000 R09: 0000000000000000
[ 16.743424] R10: 0000000000000000 R11: 0000000000000246 R12: 0000000000000000
[ 16.743424] R13: 00007ffcfa214abc R14: 00000000000925ae0 R15: 0000000000000000
[ 16.743426] ---[ end trace 8bf4d15315659a3e ]---
[ 16.743431] NVRM: GPU 0000:00:0a.0: Failed to enable MSI; falling back to PCIe virtual-
wire interrupts.

```

Make sure to run `mkinitrd` and reboot after disabling the Nouveau driver. Refer to [Section 5.1, “Prepare the VM Guest”](#).

9.6 Filing an NVIDIA vGPU bug

While filing an NVIDIA vGPU-related bug report to us, please attach the vGPU configuration data `nvidia-bug-report.log.gz` collected by the `nvidia-bug-report.sh` utility. Make sure you cover both VM Host Server and VM Guest.

9.7 Configuring a License Server

Refer to <https://docs.nvidia.com/grid/lis/latest/grid-license-server-user-guide/index.html>.

10 For more information

NVIDIA has an extensive documentation on vGPU. Refer to <https://docs.nvidia.com/grid/latest/grid-vgpu-user-guide/index.html> for details.

11 NVIDIA virtual GPU background

11.1 NVIDIA GPU architectures

There are two types of GPU architectures:

Time-sliced vGPU architecture

Introduced on GPUs that are based on the NVIDIA Ampere GPU architecture. Only Ampere GPU cards can support MIG-backed vGPU.

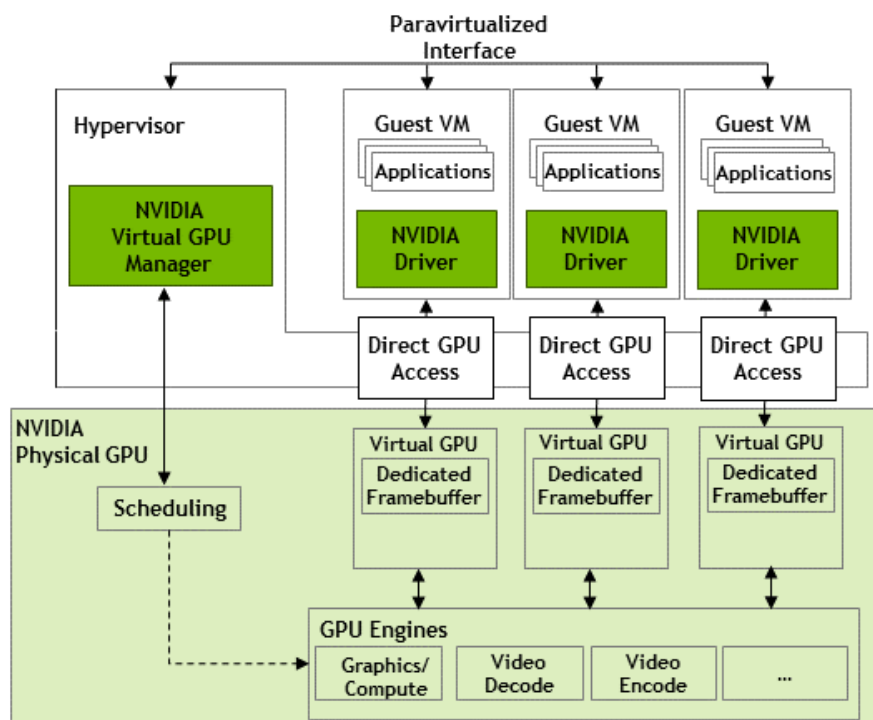


FIGURE 1: TIME-SLICED ARCHITECTURE (SOURCE: [HTTPS://DOCS.NVIDIA.COM/GRID/LATEST/GRID-VGPU-USER-GUIDE/INDEX.HTML](https://docs.nvidia.com/grid/latest/grid-vgpu-user-guide/index.html))

Multi-Instance GPU (MIG) vGPU architecture

All GPU cards support time-sliced vGPU. To do so, Ampere GPU cards use the Single Root I/O Virtualization (SR-IOV) mechanism, while Volta and the earlier GPU cards use the mediated device mechanism. Volta and the earlier architecture are based on mediated device mechanism. These two mechanisms are transparent to a VM. However, they need different configurations from the host side.

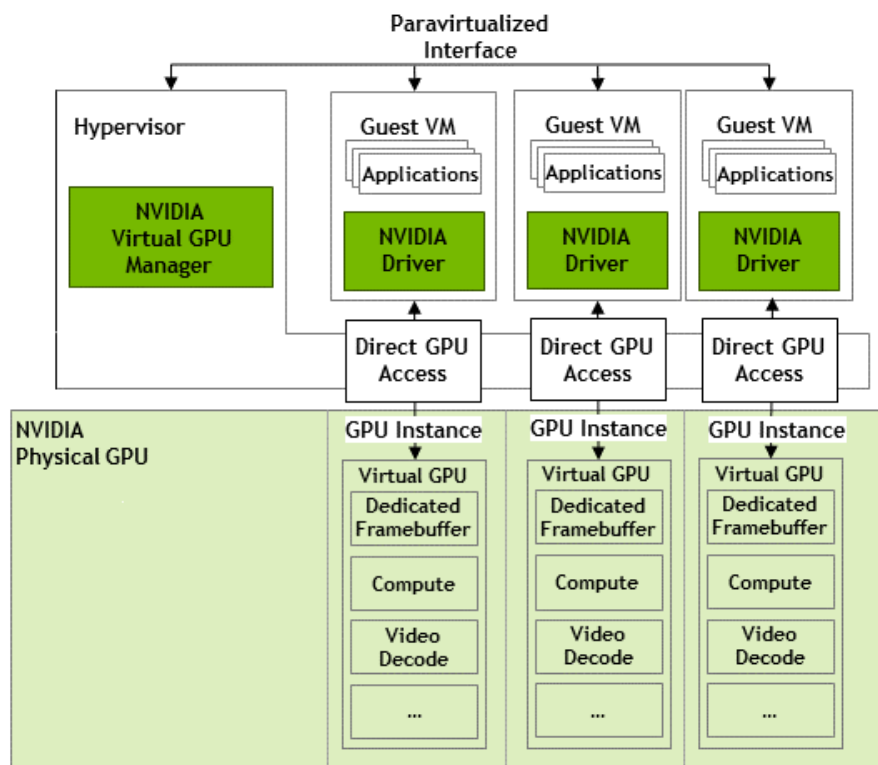


FIGURE 2: MIG-BACKED ARCHITECTURE (SOURCE: [HTTPS://DOCS.NVIDIA.COM/GRID/LATEST/GRID-VGPU-USER-GUIDE/INDEX.HTML](https://docs.nvidia.com/grid/latest/grid-vgpu-user-guide/index.html))

11.2 vGPU types

Each physical GPU can support several different types of vGPUs. vGPU types have a fixed amount of frame buffer, the number of supported display heads, and maximum resolutions. NVIDIA has four types of vGPUs: A, B, C, and Q-series. SUSE currently supports Q and C-series.

TABLE 1: VGPU TYPES

vGPU series	Optimal workload
Q-series	Virtual workstations for creative and technical professionals who require the performance and features of the NVIDIA Quadro technology.
C-series	Compute-intensive server workloads, for example, artificial intelligence (AI), deep learning, or high-performance computing (HPC).
B-series	Virtual desktops for business professionals and knowledge workers.

vGPU series	Optimal workload
A-series	Application streaming or session-based solutions for virtual applications users.

11.3 Valid vGPU configurations on a single GPU

11.3.1 Time-sliced vGPU configurations

For time-sliced vGPUs, all vGPUs types must be the same:

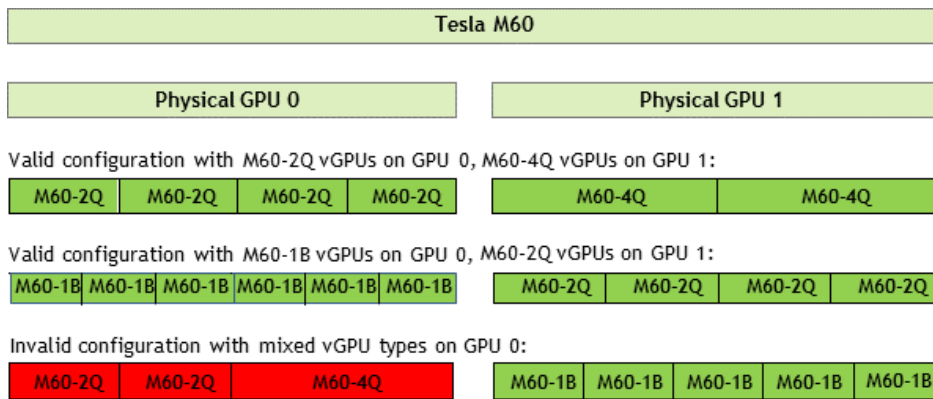


FIGURE 3: EXAMPLE TIME-SLICED VGPU CONFIGURATIONS ON NVIDIA TESLA M60 (SOURCE: [HTTPS://DOCS.NVIDIA.COM/GRID/LATEST/GRID-VGPU-USER-GUIDE/INDEX.HTML](https://docs.nvidia.com/grid/latest/grid-vgpu-user-guide/index.html))

11.3.2 MIG-backed vGPU configurations

For MIG-backed vGPUs, vGPUs can be both homogeneous and mixed-type:

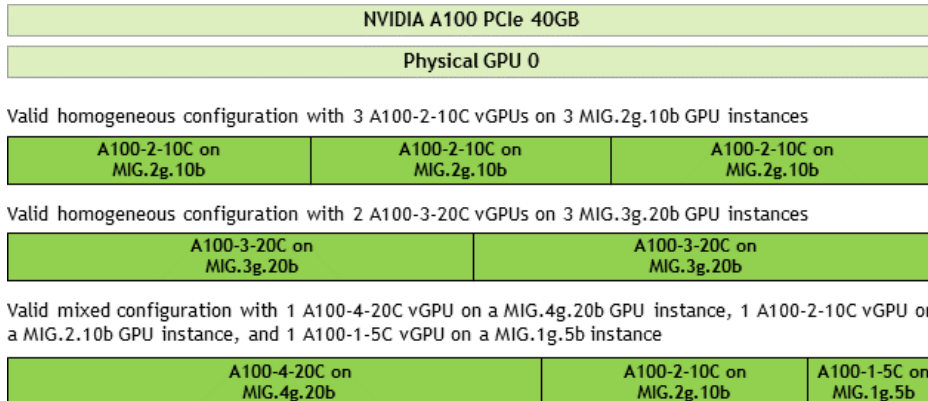


FIGURE 4: EXAMPLE MIG-BACKED vGPU CONFIGURATIONS ON NVIDIA A100 PCIe 40GB (SOURCE: [HTTPS://DOCS.NVIDIA.COM/GRID/LATEST/GRID-VGPU-USER-GUIDE/INDEX.HTML](https://docs.nvidia.com/grid/latest/grid-vgpu-user-guide/index.html))

GNU Free Documentation License

Copyright (C) 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or non-commercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties--for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <https://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

```
Copyright (c) YEAR YOUR NAME.
Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.2
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.
A copy of the license is included in the section entitled "GNU
Free Documentation License".
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the "with...Texts." line with this:

```
with the Invariant Sections being LIST THEIR TITLES, with the
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.