GRID SOFTWARE FOR RED HAT ENTERPRISE LINUX WITH KVM VERSION 367.132/370.39

RN-08687-001 _v4.9 | September 2019

Release Notes

TABLE OF CONTENTS

Chapter 1. RELEASE NOTES

These *Release Notes* summarize current status, information on validated platforms, and known issues with NVIDIA GRID™ software and hardware on Red Hat Enterprise Linux with KVM.

The most current version of the documentation for this release of NVIDIA GRID Software can be found online at GRID 4.9 Software Documentation.

This release includes the following software:

- NVIDIA Windows drivers for vGPU version 370.39
- NVIDIA Linux drivers for vGPU version 367.133

Updates in this release:

- Miscellaneous bug fixes
- Security updates

Chapter 2. VALIDATED PLATFORMS

This release of NVIDIA GRID software provides support for several NVIDIA GPUs on validated server hardware platforms, Red Hat Enterprise Linux with KVM hypervisor software versions, and guest operating systems.

2.1. Supported NVIDIA GPUs and Validated Server Platforms

This release of NVIDIA GRID software provides support for the following NVIDIA GPUs on Red Hat Enterprise Linux with KVM, running on validated server hardware platforms:

- GRID K1
- GRID K2
- Tesla M6
- Tesla M10
- Tesla M60

For a list of validated server platforms, refer to NVIDIA GRID Certified Servers.

2.2. Hypervisor Software Releases

This release supports **only** the hypervisor software release listed in the table.

If a specific release, even an update release, is not listed, it's **not** supported.

Software	Releases Supported
Red Hat Enterprise Linux with KVM	7.0-7.4

2.3. Guest OS Support

NVIDIA GRID software supports several Windows releases and Linux distributions as a guest OS using GPU pass-through.

Use only a guest OS release that is listed as supported by NVIDIA GRID software with your virtualization software. To be listed as supported, a guest OS release must be supported not only by NVIDIA GRID software, but also by your virtualization software. NVIDIA **cannot** support guest OS releases that your virtualization software does not support.

2.3.1. Windows Guest OS Support

Red Hat Enterprise Linux with KVM supports Windows guest operating systems only under specific subscription programs. For details, see Certified guest operating systems for Red Hat Enterprise Linux with KVM.

NVIDIA GRID software supports **only** the following Windows releases as a guest OS on Red Hat Enterprise Linux with KVM:

If a specific release, even an update release, is not listed, it's **not** supported.

- Windows Server 2012 R2
- Windows Server 2008 R2
- Windows 10 RTM (1507), November Update (1511), Anniversary Update (1607), Creators Update (1703) (32/64-bit)
- Windows 8.1 (32/64-bit)
- Windows 7 (32/64-bit)

2.3.2. Linux Guest OS Support

NVIDIA GRID software supports only the following 64-bit Linux distributions as a guest OS **only** on supported Tesla GPUs on Red Hat Enterprise Linux with KVM:

If a specific release, even an update release, is not listed, it's **not** supported.

- Red Hat Enterprise Linux 7.0-7.4
- CentOS 7.0-7.4
- Red Hat Enterprise Linux 6.6
- CentOS 6.6

GRID K1 and GRID K2 do not support vGPU on a Linux guest OS.

Chapter 3. SECURITY UPDATES

3.1. Restricting Access to GPU Performance Counters

The NVIDIA graphics driver contains a vulnerability (CVE-2018-6260) that may allow access to application data processed on the GPU through a side channel exposed by the GPU performance counters. To address this vulnerability, update the driver and restrict access to GPU performance counters to allow access only by administrator users and users who need to use CUDA profiling tools.

The GPU performance counters that are affected by this vulnerability are the hardware performance monitors used by the CUDA profiling tools such as CUPTI, Nsight Graphics, and Nsight Compute. These performance counters are exposed on the hypervisor host and in guest VMs only as follows:

- On the hypervisor host, they are always exposed. However, the Virtual GPU Manager does not access these performance counters and, therefore, is not affected.
- In Windows and Linux guest VMs, they are exposed **only** in VMs configured for GPU pass through. They are not exposed in VMs configured for NVIDIA vGPU.

3.1.1. Windows: Restricting Access to GPU Performance Counters for One User by Using **NVIDIA Control Panel**

Perform this task from the guest VM to which the GPU is passed through.

Ensure that you are running **NVIDIA Control Panel** version 8.1.950.

1. Open **NVIDIA Control Panel**:
   - Right-click on the Windows desktop and select **NVIDIA Control Panel** from the menu.
   - Open **Windows Control Panel** and double-click the **NVIDIA Control Panel** icon.

2. In **NVIDIA Control Panel**, select the **Manage GPU Performance Counters** task in the **Developer** section of the navigation pane.
3. Complete the task by following the instructions in the **Manage GPU Performance Counters** > **Developer** topic in the **NVIDIA Control Panel** help.

3.1.2. Windows: Restricting Access to GPU Performance Counters Across an Enterprise by Using a Registry Key

You can use a registry key to restrict access to GPU Performance Counters for all users who log in to a Windows guest VM. By incorporating the registry key information into a script, you can automate the setting of this registry for all Windows guest VMs across your enterprise.

Perform this task from the guest VM to which the GPU is passed through.

Caution Only **enterprise administrators** should perform this task. Changes to the Windows registry must be made with care and system instability can result if registry keys are incorrectly set.

1. Set the `RmProfilingAdminOnly` Windows registry key to 1.

```
[HKEY_LOCAL_MACHINE\SOFTWARE\NVIDIA Corporation\Global\NVTweak]
Value: "RmProfilingAdminOnly"
Type: DWORD
Data: 00000001
```

   The data value 1 restricts access, and the data value 0 allows access, to application data processed on the GPU through a side channel exposed by the GPU performance counters.
2. Restart the VM.

3.1.3. Linux Guest VMs and Hypervisor Host: Restricting Access to GPU Performance Counters

On systems where unprivileged users don't need to use GPU performance counters, restrict access to these counters to system administrators, namely users with the `CAP_SYS_ADMIN` capability set. By default, the GPU performance counters are not restricted to users with the `CAP_SYS_ADMIN` capability.

Perform this task from the guest VM to which the GPU is passed through or from your hypervisor host machine.

In Linux guest VMs, this task requires `sudo` privileges. On your hypervisor host machine, this task must be performed as the root user on the machine.

1. Log in to the guest VM or open a command shell on your hypervisor host machine.
2. Set the kernel module parameter `NVreg_RestrictProfilingToAdminUsers` to 1 by adding this parameter to the `/etc/modprobe.d/nvidia.conf` file.

- If you are setting only this parameter, add an entry for it to the `/etc/modprobe.d/nvidia.conf` file as follows:

```
options nvidia
 NVreg_RegistryDwords="NVreg_RestrictProfilingToAdminUsers=1"
```

- If you are setting multiple parameters, set them in a single entry as in the following example:

```
options nvidia NVreg_RegistryDwords="RmPVMRL=0x0"
 "NVreg_RestrictProfilingToAdminUsers=1"
```

If the `/etc/modprobe.d/nvidia.conf` file does not already exist, create it.

3. Restart the VM or reboot your hypervisor host machine.

Notice

HDMI

OpenCL

Trademarks

Copyright