



# What's New in Virtual GPU Software R430 for All Supported Hypervisors

Release Notes

# Table of Contents

- Chapter 1. Updates by Release..... 1
  - 1.1. Updates in Release 9.0..... 1
  - 1.2. Updates in Release 9.1..... 2
  - 1.3. Updates in Release 9.2..... 2
  - 1.4. Updates in Release 9.3..... 3
  - 1.5. Updates in Release 9.4..... 3

---

# Chapter 1. Updates by Release

Updates for each release in this release family of NVIDIA vGPU software may include new features, introduction of hardware and software support, and withdrawal of hardware and software support.

## 1.1. Updates in Release 9.0

### New Features in Release 9.0

- ▶ NVIDIA Virtual Compute Server (VCS) vGPUs for artificial intelligence, deep learning, and high-performance computing workloads
- ▶ Support for multiple vGPUs in a single VM on VMware vSphere (requires release 6.7 Update 3)
- ▶ Error correcting code (ECC) memory support
- ▶ Page retirement support
- ▶ Peer-to-peer CUDA transfers over NVLink support
- ▶ Configurable time slices for equal share schedulers and fixed share schedulers
- ▶ New configuration parameter to specify host ID of a licensed client
- ▶ Miscellaneous bug fixes

### Hardware and Software Support Introduced in Release 9.0

- ▶ Support for the Tesla T4 GPU on Nutanix AHV
- ▶ Support for Nutanix AHV 5.11
- ▶ Support for Windows 10 May 2019 Update (1903) as a guest OS
- ▶ Support for Ubuntu 18.04 LTS as a guest OS on Citrix Hypervisor 7.1 cumulative update 2 and Microsoft Windows Server
- ▶ Support for VMware Horizon 7.9

### Feature Support Withdrawn in Release 9.0

- ▶ VMware vSphere ESXi 6.0 is no longer supported.

## 1.2. Updates in Release 9.1

### New Features in Release 9.1

- ▶ Support for NVIDIA Virtual Compute Server vGPUs on the following GPUs:
  - ▶ Quadro RTX 6000
  - ▶ Quadro RTX 8000
- ▶ Security updates
- ▶ Miscellaneous bug fixes

### Hardware and Software Support Introduced in Release 9.1

- ▶ Support for CentOS 7.7 as a guest OS
- ▶ Support for Red Hat Enterprise Linux 7.7 as a guest OS
- ▶ Support for the following versions of Citrix Virtual Apps and Desktops:
  - ▶ Version 7 1906
  - ▶ Version 7 1909

## 1.3. Updates in Release 9.2

### New Features in Release 9.2

- ▶ Miscellaneous bug fixes
- ▶ Security updates
- ▶ Limitation on the maximum number of NVIDIA Virtual Compute Server vGPUs to eight vGPUs per physical GPU, irrespective of the available hardware resources of the physical GPU
- ▶ Inclusion of two Virtual GPU Manager RPM packages for Red Hat Enterprise Linux with KVM 8.0:
  - ▶ One package for the GA Linux KVM kernel
  - ▶ One package for the updated z-stream kernel

### Hardware and Software Support Introduced in Release 9.2

- ▶ Support for VMware Horizon 7.10

## 1.4. Updates in Release 9.3

### New Features in Release 9.3

- ▶ Miscellaneous bug fixes
- ▶ Security updates (see [Security Bulletin: NVIDIA GPU Display Driver - February 2020](#))

### Hardware and Software Support Introduced in Release 9.3

- ▶ Support for Citrix Virtual Apps and Desktops version 7 1912
- ▶ Support for VMware Horizon 7.12 and 7.11

### Feature Support Withdrawn in Release 9.3

- ▶ Citrix Hypervisor 7.6 is no longer supported.
- ▶ Red Hat Enterprise Linux with KVM hypervisor releases 7.0-7.4 are no longer supported.
- ▶ The following OS releases are no longer supported as a guest OS:
  - ▶ Windows Server 2008 R2
  - ▶ Red Hat Enterprise Linux 7.0-7.4
  - ▶ CentOS 7.0-7.4

## 1.5. Updates in Release 9.4

### New Features in Release 9.4

- ▶ Miscellaneous bug fixes
- ▶ Security updates - see [Security Bulletin: NVIDIA GPU Display Driver - June 2020](#)

### Hardware and Software Support Introduced in Release 9.4

- ▶ Support for Nutanix AHV 5.15
- ▶ Support for Red Hat Enterprise Linux with KVM hypervisor release 7.8
- ▶ Support for the following OS releases as a guest OS:
  - ▶ Red Hat Enterprise Linux 7.8
  - ▶ CentOS 7.8

### Feature Support Withdrawn in Release 9.4

- ▶ Red Hat Enterprise Linux with KVM hypervisor 7.5 is no longer supported.

- ▶ Nutanix AHV 5.9 and 5.8 are no longer supported
- ▶ The following OS releases are no longer supported as a guest OS:
  - ▶ Red Hat Enterprise Linux 7.5
  - ▶ CentOS 7.5

## Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

## VESA DisplayPort

DisplayPort and DisplayPort Compliance Logo, DisplayPort Compliance Logo for Dual-mode Sources, and DisplayPort Compliance Logo for Active Cables are trademarks owned by the Video Electronics Standards Association in the United States and other countries.

## HDMI

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

## OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

## Trademarks

NVIDIA, the NVIDIA logo, NVIDIA GRID, NVIDIA GRID vGPU, NVIDIA Maxwell, NVIDIA Pascal, NVIDIA Turing, NVIDIA Volta, Quadro, and Tesla are trademarks or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

## Copyright

© 2013-2020 NVIDIA Corporation. All rights reserved.