



What's New in Virtual GPU Software R535 for All Supported Hypervisors

Release Notes

Table of Contents

Chapter 1. Updates by Release.....	1
1.1. Updates in Release 16.6.....	1
1.2. Updates in Release 16.5.....	2
1.3. Updates in Release 16.4.....	2
1.4. Updates in Release 16.3.....	2
1.5. Updates in Release 16.2.....	3
1.6. Updates in Release 16.1.....	3
1.7. Updates in Release 16.0.....	4

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 16.6

New Features in Release 16.6

- ▶ vGPU migration support on the Red Hat Enterprise Linux with KVM
- ▶ Security updates - see *Security Bulletin: NVIDIA GPU Display Driver - June 2024*, which is posted shortly after the release date of this software and is listed on the [NVIDIA Product Security](#) page
- ▶ Miscellaneous bug fixes

Hardware and Software Support Introduced in Release 16.6

- ▶ Newly supported hypervisor software:
 - ▶ Red Hat Enterprise Linux with KVM 9.4 and 8.10
- ▶ Newly supported guest OS releases:
 - ▶ Debian 12
 - ▶ Red Hat Enterprise Linux 9.4 and 8.10
 - ▶ Ubuntu 24.04



Note: This release introduces support for Ubuntu 24.04 **only** as a guest OS. This release does **not** introduce support for Ubuntu 24.04 hypervisor.

- ▶ Newly supported remoting solutions:
 - ▶ Citrix Virtual Apps and Desktops version 7 2402

Feature Support Withdrawn in Release 16.6

- ▶ Hypervisor software no longer supported:

- ▶ Red Hat Enterprise Linux with KVM 9.3, 9.0, 8.9, and 8.6
- ▶ Guest OSes no longer supported:
 - ▶ Debian 10
 - ▶ Red Hat Enterprise Linux 9.3, 9.0, 8.9, and 8.6

1.2. Updates in Release 16.5

NVIDIA vGPU software 16.5 resolves an issue that affects graphics cards that are supported only by NVIDIA AI Enterprise.

1.3. Updates in Release 16.4

New Features in Release 16.4

- ▶ Security updates - see *Security Bulletin: NVIDIA GPU Display Driver - February 2024*, which is posted shortly after the release date of this software and is listed on the [NVIDIA Product Security](#) page
- ▶ Miscellaneous bug fixes

Newly Supported Hardware and Software in Release 16.4

- ▶ Newly supported hypervisor software:
 - ▶ Red Hat Enterprise Linux with KVM hypervisor 9.3
- ▶ Newly supported guest OS releases:
 - ▶ Microsoft Windows 11 23H2
 - ▶ Red Hat Enterprise Linux 9.3
- ▶ Newly supported remoting solutions:
 - ▶ VMware Horizon 2312 (8.12)

1.4. Updates in Release 16.3

New Features in Release 16.3

- ▶ Miscellaneous bug fixes

Hardware and Software Support Introduced in Release 16.3

- ▶ Newly supported graphics cards:

- ▶ NVIDIA L2
- ▶ NVIDIA L20
- ▶ Newly supported hypervisor software:
 - ▶ Red Hat Enterprise Linux with KVM hypervisor 8.9
- ▶ Newly supported guest OS releases:
 - ▶ Red Hat Enterprise Linux 8.9
- ▶ Newly supported remoting solutions:
 - ▶ Citrix Virtual Apps and Desktops version 7 2311

1.5. Updates in Release 16.2

New Features in Release 16.2

- ▶ Security updates - see *Security Bulletin: NVIDIA GPU Display Driver - October 2023*, which is posted shortly after the release date of this software and is listed on the [NVIDIA Product Security](#) page
- ▶ Miscellaneous bug fixes

Hardware and Software Support Introduced in Release 16.2

- ▶ Newly supported hypervisor software:
 - ▶ Microsoft Azure Stack HCI 23H2 preview
- ▶ Newly supported remoting solutions:
 - ▶ Citrix Virtual Apps and Desktops version 7 2308
 - ▶ VMware Horizon 2309 (8.11)

1.6. Updates in Release 16.1

New Features in Release 16.1

- ▶ New options in the NVML API and the `nvidia-smi` command for getting the scheduling behavior of time-sliced vGPUs
- ▶ Miscellaneous bug fixes

Hardware and Software Support Introduced in Release 16.1

- ▶ Support for the for the following GPUs:
 - ▶ NVIDIA L40S

- ▶ NVIDIA RTX 5000 Ada
- ▶ Support for XenServer 8 hypervisor
- ▶ Support for Microsoft Windows Server 2022 as a guest OS on Red Hat Enterprise Linux with KVM hypervisor
- ▶ Support for SUSE Linux Enterprise Server 12 SP5 as a guest OS on Microsoft Windows Server and VMware vSphere

1.7. Updates in Release 16.0

New Features in Release 16.0

- ▶ Support for the NVIDIA L4 and NVIDIA L40 graphics cards in DDA mode only on Microsoft Azure Stack HCI
- ▶ Support for 4K displays with an aspect ratio of 16:10
- ▶ Options in the NVML API and the `nvidia-smi` command for controlling the scheduling behavior of time-sliced vGPUs
- ▶ Assignment of multiple fractional vGPUs to a single VM on Microsoft Azure Stack HCI
A fractional vGPU is allocated only a fraction of the physical GPU's frame buffer.
- ▶ Security updates - see *Security Bulletin: NVIDIA GPU Display Driver - June 2023*, which is posted shortly after the release date of this software and is listed on the [NVIDIA Product Security](#) page
- ▶ Miscellaneous bug fixes

Newly Supported Hardware and Software in Release 16.0

- ▶ Newly supported hypervisor software:
 - ▶ Red Hat Enterprise Linux with KVM 9.2 and 8.8
- ▶ Newly supported guest OSes:
 - ▶ Red Hat Enterprise Linux 9.2 and 8.8
- ▶ Newly supported remoting solutions:
 - ▶ Citrix Virtual Apps and Desktops version 7 2305
 - ▶ VMware Horizon 2306 (8.10) and 2303 (8.9)

Feature Support Withdrawn in Release 16.0

- ▶ Product functionality no longer supported:

- ▶ NVIDIA Virtual Compute Server (vCS)



Note: As a result of this change, C-series vGPU types are no longer available.

Instead, vCS is supported with NVIDIA AI Enterprise.

- ▶ Graphics cards no longer supported:
 - ▶ Graphics cards that support only C-series vGPUs, namely:
 - ▶ NVIDIA H800 PCIe 80GB
 - ▶ NVIDIA H100 PCIe 80GB
 - ▶ NVIDIA A800 PCIe 80GB
 - ▶ NVIDIA A800 PCIe 80GB liquid cooled
 - ▶ NVIDIA A800 HGX 80GB
 - ▶ NVIDIA A100 PCIe 80GB
 - ▶ NVIDIA A100 PCIe 80GB liquid cooled
 - ▶ NVIDIA A100X
 - ▶ NVIDIA A100 HGX 80GB
 - ▶ NVIDIA A100 PCIe 40GB
 - ▶ NVIDIA A100 HGX 40GB
 - ▶ NVIDIA A30
 - ▶ NVIDIA A30X

Instead, these graphics cards are supported with NVIDIA AI Enterprise.

- ▶ Hypervisor software no longer supported:
 - ▶ Red Hat Enterprise Linux with KVM 9.1, 8.7, and 8.4
 - ▶ Ubuntu 18.04
- ▶ Guest OSes no longer supported:
 - ▶ Red Hat Enterprise Linux 9.1, 8.7, and 8.4
 - ▶ Ubuntu 18.04

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, GPUDirect, 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

© 2013-2024 NVIDIA Corporation & affiliates. All rights reserved.

