



What's New in Virtual GPU Software R510 for All Supported Hypervisors

Release Notes

Table of Contents

- Chapter 1. Updates by Release..... 1
 - 1.1. Updates in Release 14.4..... 1
 - 1.2. Updates in Release 14.3..... 1
 - 1.3. Updates in Release 14.2..... 2
 - 1.4. Updates in Release 14.1..... 2
 - 1.5. Updates in Release 14.0..... 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 14.4

New Features in Release 14.4

- ▶ Security updates - see *Security Bulletin: NVIDIA GPU Display Driver - November 2022*, which is updated shortly after the release date of this software and is listed on the [NVIDIA Product Security](#) page

1.2. Updates in Release 14.3

New Features in Release 14.3

- ▶ Security updates - see *Security Bulletin: NVIDIA GPU Display Driver - November 2022*, which is posted shortly after the release date of this software and is listed on the [NVIDIA Product Security](#) page
- ▶ Support for non-transparent local proxy servers when NVIDIA vGPU software is served licenses by a Cloud License Service (CLS) instance
- ▶ Miscellaneous bug fixes

Hardware and Software Support Introduced in Release 14.3

- ▶ Support for Red Hat Enterprise Linux with KVM hypervisor 8.7
- ▶ Support for Red Hat Enterprise Linux 8.7 as a guest OS
- ▶ Support for Citrix Virtual Apps and Desktops version 7 2209
- ▶ Support for VMware Horizon 2209 (8.7)

1.3. Updates in Release 14.2

New Features in Release 14.2

- ▶ Security updates - see *Security Bulletin: NVIDIA GPU Display Driver - August 2022*, 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 14.2

- ▶ Support for Citrix Virtual Apps and Desktops version 7 2206
- ▶ Support for Ubuntu 22.04 LTS on VMware vSphere Hypervisor (ESXi)
- ▶ Support for VMware Horizon 2206 (8.6)

Feature Support Withdrawn in Release 14.2

- ▶ Citrix XenServer 7.1 is no longer supported.
- ▶ All versions of Microsoft Windows Server 2016 with Hyper-V role are no longer supported as a hypervisor.
- ▶ VMware vSphere Hypervisor (ESXi) 6.7 and 6.5 are no longer supported.
- ▶ All versions of Microsoft Windows Server 2016 are no longer supported as a guest OS.

1.4. Updates in Release 14.1

New Features in Release 14.1

- ▶ Support for GPUDirect[®] technology on the Mellanox Connect-X[®] 7 SmartNIC network interface card
- ▶ Security updates - see *Security Bulletin: NVIDIA GPU Display Driver - May 2022*, 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 14.1

- ▶ Support for Red Hat Enterprise Linux with KVM hypervisor 9.0 and 8.6
- ▶ Support for Ubuntu 22.04 LTS as a hypervisor
- ▶ Support for Red Hat Enterprise Linux 9.0 as a guest OS on the Red Hat Enterprise Linux with KVM hypervisor

- ▶ Support for Red Hat Enterprise Linux 8.6 as a guest OS on the Citrix Hypervisor, Microsoft Windows Server, Red Hat Enterprise Linux with KVM, and VMware vSphere hypervisors
- ▶ Support for Ubuntu 22.04 LTS as a guest OS on the Ubuntu hypervisor
- ▶ Support for Citrix Virtual Apps and Desktops version 7 2203
- ▶ Support for VMware Horizon 2203 (8.5)

Feature Support Withdrawn in Release 14.1

- ▶ Red Hat Enterprise Linux with KVM hypervisor 8.5 is no longer supported.
- ▶ Red Hat Enterprise Linux 8.5 and 8.2 are no longer supported as a guest OS.

1.5. Updates in Release 14.0

New Features in Release 14.0

- ▶ Support for NVIDIA GPUDirect[®] Storage technology on the generic Linux with KVM, Red Hat Enterprise Linux with KVM, and Ubuntu hypervisors
- ▶ Support for GPUDirect technology on all C-series vGPUs on GPUs that support SR-IOV on the Linux with KVM, Red Hat Enterprise Linux with KVM, and Ubuntu hypervisors
- ▶ Support for a mixture of time-sliced vGPUs of the same frame buffer size on the same GPU on the Linux with KVM and Red Hat Enterprise Linux with KVM hypervisors
- ▶ Support for Tesla Compute Cluster (TCC) mode for Q-series vGPUs on Windows guest VMs on the Citrix Hypervisor, Red Hat Enterprise Linux with KVM, and VMware vSphere hypervisors
- ▶ Support for GPU System Processor (GSP) in GPU pass through and bare-metal configurations on Linux with vCS



Note: If you are using a product other than vCS, you must disable GSP as explained in [Virtual GPU Software User Guide](#).

- ▶ Enhanced NVIDIA CUDA Toolkit support:
 - ▶ NVIDIA CUDA Toolkit profilers can be enabled when unified memory is enabled.
 - ▶ Nsight Systems GPU context switch trace is supported.
- ▶ Enhancements to the NVIDIA Management Library (NVML) to determine whether a vGPU type supports GPUDirect technology and peer-to-peer CUDA transfers over NVLink
- ▶ Addition of RPM and Debian packages for the NVIDIA vGPU software graphics drivers for Linux
- ▶ Security updates - see *Security Bulletin: NVIDIA GPU Display Driver - February 2022*, 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 14.0

- ▶ Support for the following GPUs:
 - ▶ NVIDIA A2
 - ▶ NVIDIA A30X
 - ▶ NVIDIA A100X
 - ▶ NVIDIA RTX A5500
- ▶ Support for Red Hat Enterprise Linux with KVM hypervisor 8.5
- ▶ Support for Red Hat Enterprise Linux 8.5 as a guest OS on the Citrix Hypervisor, Microsoft Windows Server, Red Hat Enterprise Linux with KVM, and VMware vSphere hypervisors
- ▶ Support for Red Hat Enterprise Linux 8.4 and 8.2 as a guest OS on Microsoft Windows Server
- ▶ Support for Debian 10 as a guest OS on VMware vSphere
- ▶ Support for Ubuntu 20.04 LTS as a guest OS on Microsoft Windows Server
- ▶ Support for Citrix Virtual Apps and Desktops version 7 2112
- ▶ Support for VMware Horizon 2111 (8.4)

Feature Support Withdrawn in Release 14.0

- ▶ Red Hat Enterprise Linux with KVM hypervisor 8.1, 7.8, and 7.7 are no longer supported.
- ▶ Red Hat Enterprise Linux 8.1 is no longer supported as a guest OS.
- ▶ Red Hat Enterprise Linux 7.8 and 7.7 are no longer supported as a guest OS.
- ▶ Windows Server 2012 R2 is no longer supported as a guest OS.

Features Deprecated in Release 14.0

The following table lists features that are deprecated in this release of NVIDIA vGPU software. Although the features remain available in this release, they might be withdrawn in a future release. In preparation for the possible removal of these features, use the preferred alternative listed in the table.

Deprecated Feature	Preferred Alternative	Additional Information
Legacy NVIDIA vGPU software license server	NVIDIA License System	NVIDIA Virtual GPU Software License Server End of Life Notice

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

© 2013-2022 NVIDIA Corporation & affiliates. All rights reserved.

