What's New in Virtual GPU Software R580 for All Supported Hypervisors

Release Notes

RN-09409-001 _v19.0 through 19.4 | January 2026

Table of Contents

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 19.4

New Features in Release 19.4

- Security updates

  For the latest NVIDIA vulnerability disclosure information, visit the NVIDIA Product Security page. Security updates for NVIDIA vGPU software are described the NVIDIA GPU Display Driver security bulletins.
- Miscellaneous bug fixes

Hardware and Software Support Introduced in Release 19.4

- Newly supported hypervisor software:
  - Red Hat Enterprise Linux with KVM 10.1
  - Red Hat Enterprise Linux with KVM 9.7
- Newly supported guest OS releases:
  - Red Hat Enterprise Linux 10.1

    Note: Because NVIDIA vGPU software does not support the Wayland display server protocol, Red Hat Enterprise Linux 10.1 is supported only in console and command-line interface (CLI) mode.
  - Red Hat Enterprise Linux 9.7
  - Ubuntu 24.04 LTS on XenServer 8.4
  - Windows 11 25H2
- Newly supported remoting solutions:
  - Citrix Virtual Apps and Desktops version 7 2511

- Omnissa Horizon 8 2512

1.2. Updates in Release 19.3

New Features in Release 19.3

- Support for Windows Enterprise multi-session with Windows 11 guest OS releases on the Microsoft Azure Local hypervisor
- Miscellaneous bug fixes

1.3. Updates in Release 19.2

Note: Starting with NVIDIA vGPU software 19.2, support for NVIDIA vGPU software on VMware vSphere Hypervisor (ESXi) 8.0 requires VMware vSphere Hypervisor (ESXi) 8.0u3 P06 and later updates to release 8.0 unless explicitly stated otherwise. Earlier VMware vSphere Hypervisor (ESXi) 8.0 builds are **not** supported.

New Features in Release 19.2

- Security updates

  For the latest NVIDIA vulnerability disclosure information, visit the NVIDIA Product Security page. Security updates for NVIDIA vGPU software are described the NVIDIA GPU Display Driver security bulletins.
- Miscellaneous bug fixes

Hardware and Software Support Introduced in Release 19.2

- Newly supported graphics cards:
  - NVIDIA RTX PRO 6000 Blackwell Server Edition on VMware vSphere ESXi hypervisor 9.0.1.0
  - NVIDIA RTX PRO 6000D on the Red Hat Enterprise Linux with KVM and Ubuntu hypervisors
- Newly supported guest OS releases:
  - Azure Linux 3
  - Microsoft Windows Server 2025 on XenServer

1.4. Updates in Release 19.1

New Features in Release 19.1

- New B-series vGPU profiles with 3 GB of frame buffer on supported GPUs based on the NVIDIA Ampere GPU architectures
- vGPU hibernation support for GPUs based on the NVIDIA Ada Lovelace GPU architecture on Microsoft Azure Local
- Miscellaneous bug fixes

Hardware and Software Support Introduced in Release 19.1

- Newly supported graphics cards:
  - NVIDIA RTX PRO 6000 Blackwell Server Edition on the XenServer, Microsoft Azure Local, and Microsoft Windows Server hypervisors
- Newly supported remoting solutions:
  - Citrix Virtual Apps and Desktops version 7 2507
  - Omnissa Horizon 8 2506

1.5. Updates in Release 19.0

New Features in Release 19.0

- Support for Multi-Instance GPU (MIG)-backed vGPUs for graphics on GPUs that support MIG
- Support for time-sliced, MIG-backed vGPUs within a GPU instance on a MIG-enabled GPU
- New B-series vGPU profiles with 3 GB of frame buffer on supported GPUs based on the NVIDIA Ada Lovelace and NVIDIA Blackwell GPU architectures
- Support for Virtualization Based Security (VBS) in Windows 11 guest VMs
- Miscellaneous bug fixes

Newly Supported Hardware and Software in Release 19.0

- Newly supported graphics cards:
  - NVIDIA RTX PRO 6000 Blackwell Server Edition on the Red Hat Enterprise Linux with KVM, Ubuntu, and VMware vSphere hypervisors
- Newly supported guest OS releases:

- Ubuntu 24.04 LTS and Ubuntu 22.04 LTS on Microsoft Windows Server 2025 in GPU-P mode only

Features Deprecated in Release 19.0

NVIDIA vGPU software 19 is the last release branch to support the following graphics cards:

- Tesla M10
- Tesla V100 SXM2
- Tesla V100 SXM2 32GB
- Tesla V100 PCIe
- Tesla V100 PCIe 32GB
- Tesla V100S PCIe 32GB
- Tesla V100 FHHL
- Quadro RTX 6000
- Quadro RTX 6000 passive
- Quadro RTX 8000
- Quadro RTX 8000 passive

Disabling strict round robin policy is deprecated and NVIDIA vGPU software 19 is the last release branch to support it. Support for this feature is planned to be removed in the next major release of NVIDIA vGPU software.

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 NONINFRINGEMENT, 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

NVIDIA Corporation | 2788 San Tomas Expressway, Santa Clara, CA 95051
http://www.nvidia.com