



WHAT'S NEW IN VIRTUAL GPU SOFTWARE R430 FOR ALL SUPPORTED HYPERVISORS

RN-09409-001 _v9.0 through 9.2 | November 2019

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

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 vComputeServer 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 times 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 the following OS releases as a guest OS on VMware vSphere:
 - ▶ Red Hat Enterprise Linux 8.0
 - ▶ CentOS 8.0

- ▶ 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Feature Support Withdrawn in Release 9.0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 vComputeServer 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 vComputeServer 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

Notice

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.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

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, vGPU, Pascal, 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

© 2013-2019 NVIDIA Corporation. All rights reserved.