GPU Manager for VMware vCenter

Release Notes

RN-11359-001 _v2.0.0 | February 2025

Table of Contents

Chapter 1. Supported Platforms

1.1. Supported Software Releases

NVIDIA GPU Manager for VMware vCenter is supported on specific releases of VMware vSphere Hypervisor (ESXi) and NVIDIA vGPU software.

Software	Supported Releases
VMware vCenter Server	▶ 8.0 and later updates to release 8.0 unless explicitly stated otherwise ▶ 7.0 Update 2 and later updates to release 7.0 unless explicitly stated otherwise **Note:** The base VMware vSphere Hypervisor (ESXi) 7.0 release and 7.0 Update 1 are **not** supported.
VMware vSphere with Tanzu	VMware vSphere 8.0 with supervisor cluster version v1.26.4+vmware.wcp.1-vsc0.1.6-22282210 Later updates to VMware vSphere 8.0 are also supported unless explicitly stated otherwise.
NVIDIA vGPU software	All supported releases of Virtual GPU Manager for VMware vSphere

1.2. Web Browser Requirements

The NVIDIA GPU Manager for VMware vCenter virtual appliance provides a web-based management interface to the **NVIDIA GPU Manager** application on the appliance. This management interface must be accessed through a supported web browser.

The following web browsers are supported:

- Firefox
- Google Chrome
- Microsoft Edge

- Safari

1.3. Known Product Limitations

NVIDIA GPU Manager for VMware vCenter supports **only** IPv4. IPv6 is **not** supported.

Chapter 2. Changes in this Release

Changes in Release 2.0

- Enhancements to the **Driver Download** tab of the **NVIDIA GPU Manager** page:
  - Drivers that have already been downloaded are listed.
  - Users can now filter the list of NVIDIA GPU drivers available for download.
- New, more stringent requirements for the **`vcp_cli_user`** password

  These requirements are a result of the upgrade of the base OS of the NVIDIA GPU Manager for VMware vCenter virtual appliance to Ubuntu 24.04.1. This release of Ubuntu has more stringent hardening requirements than the base OS in the previous version of the appliance.

  > **Note:** These requirements do not affect existing **`vcp_cli_user`** passwords, which remain valid.

- Security updates as listed in Security Updates
- Miscellaneous bug fixes

Chapter 3. Security Updates

Security Updates in Release 2.0

To address vulnerabilities in NVIDIA GPU Manager for VMware vCenter, new versions of the following software are included:

- NGINX (1.26.2)
- PostgreSQL (16.6)
- Python (3.12.3)
- Ubuntu OS (24.04.1)

Chapter 4. Resolved Issues

Only resolved issues that have been previously noted as known issues or had a noticeable user impact are listed. The summary and description for each resolved issue indicate the effect of the issue on NVIDIA GPU Manager for VMware vCenter **before the issue was resolved**.

4.1. Issues Resolved in Release 2.0

No resolved issues are reported in this release.

Chapter 5. Known Issues

5.1. Assignment of a static IP address fails if any properties are invalid or blank

Description

Network settings for a static IP address can be provided during the installation of the NVIDIA GPU Manager for VMware vCenter virtual appliance. However, these values cannot be validated before the VM that hosts the appliance is created. Therefore, if any field **except DNS Server 2** is blank or if any field contains invalid data, a static IP address is **not** assigned to the VM. Instead, the VM obtains network settings from a DHCP server.

Workaround

If the VM obtains network settings from a DHCP server, assign a static IP address to the VM.

1. Register the NVIDIA GPU Manager for VMware vCenter administrator user.
2. Update the network settings for the NVIDIA GPU Manager for VMware vCenter virtual appliance.

For instructions for how to perform these tasks, refer to *GPU Manager for VMware vCenter User Guide*.

Do not attempt to register the NVIDIA GPU Manager for VMware vCenter until a static IP address has been assigned to the VM that hosts the appliance.

Status

Open

Ref.

4080509

5.2. Unregistering NVIDIA GPU Manager for VMware vCenter does not remove the **NVIDIA GPU Manager drivers repository**

Description

After NVIDIA GPU Manager for VMware vCenter is unregistered, the entry for **NVIDIA GPU Manager drivers repository** is not removed from vSphere Lifecycle Management (vLCM). This issue occurs if an NVIDIA GPU Manager for VMware vCenter for which no drivers are synchronized with vLCM is unregistered.

When this issue occurs, the following error message can be seen by selecting **vSphere Client** > **Tasks** :

```
A general system error occurred: Cannot find the content of online depot
'https://nvidia-gpu-manager-appliance-ip/DriverRepo/index.xml';.
Either the depot's content has not been imported yet or has already been deleted
or the depot was migrated from an old version of
vSphere Lifecycle Manager that does not support this operation.
```

Version

This issue affects supported versions of VMware vSphere Hypervisor (ESXi) before 8.0 update 1.

Workaround

Delete the entry in the table for **NVIDIA GPU Manager drivers repository** manually as explained in Modify a Download Source on the VMware Docs site.

Status

Not an NVIDIA bug

Ref.

3974270

5.3. Unregistering NVIDIA GPU Manager for VMware vCenter does not remove downloaded NVIDIA GPU drivers from vLCM

Description

After NVIDIA GPU Manager for VMware vCenter is unregistered, the entries for downloaded NVIDIA GPU drivers are not removed from vSphere Lifecycle Management (vLCM). The downloaded drivers can still be seen by selecting **vSphere Client** > **Lifecycle Manager** > **COMPONENTS** .

Version

This issue affects supported versions of VMware vSphere Hypervisor (ESXi) before 8.0 update 1.

Workaround

Delete the entries for downloaded NVIDIA GPU drivers manually from vLCM as explained in VMware Knowledge Base Article: Resetting VMware Update Manager Database in vCenter Server Appliance 6.5/6.7/7.0 (2147284).

CAUTION:

- Resetting the Update Manager database is a destructive task. The following items are removed and must be reapplied after the reset:
  - Custom baselines (but not cluster images)
  - Custom download settings
  - Manually imported patches and ISO images
- To prevent loss of data if the reset fails, back up your data before resetting the Update Manager database as follows:
  - Create a snapshot of the virtual appliance and, if the vCenter Server is part of a Linked Mode replication setup, create a snapshot of all replicating nodes.

    For instructions, refer to Take Snapshots of a Virtual Machine.

    If the reset fails, you can restore the affected virtual appliance and replicating nodes from their snapshots as explained in Revert a Virtual Machine Snapshot.

- Note all custom configuration settings in Update Manager, for example, proxy settings, third-party download URLs, and customized baselines.

Status

Not an NVIDIA bug

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, 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