



GPU Manager for VMware vCenter

Release Notes


Table of Contents

Chapter 1. Supported Platforms.....	1
1.1. Supported Software Releases.....	1
1.2. Web Browser Requirements.....	1
1.3. Known Product Limitations.....	2
Chapter 2. Changes in this Release.....	3
Chapter 3. Security Updates.....	5
Chapter 4. Resolved Issues.....	6
4.1. Issues Resolved in Release 1.4.....	6
4.2. Issues Resolved in Release 1.3.....	6
4.3. Issues Resolved in Release 1.2.....	7
4.4. Issues Resolved in Release 1.1.....	7
Chapter 5. Known Issues.....	9
5.1. 1.3.0 Only: Static IP address cannot be set during installation of the NVIDIA GPU Manager for VMware vCenter virtual appliance.....	9
5.2. 1.0.0-1.2.0 Only: NVIDIA GPU Manager for VMware vCenter registration fails if IPv6 and IPv4 addresses are associated with the vCenter Server domain name.....	10
5.3. 1.1.0 Only: NVIDIA GPU Manager for VMware vCenter registration fails with incompatible error on vSphere 7.0.3.01700.....	11
5.4. Assignment of a static IP address fails if any properties are invalid or blank.....	12
5.5. Unregistering NVIDIA GPU Manager for VMware vCenter does not remove the NVIDIA GPU Manager drivers repository.....	13
5.6. Unregistering NVIDIA GPU Manager for VMware vCenter does not remove downloaded NVIDIA GPU drivers from vLCM.....	14
5.7. Since 1.2.0: System Text on the Proxy Server Settings page is not localized.....	15
5.8. 1.0.0, 1.1.0 Only: NVIDIA GPU Manager does not support communication through a proxy server.....	15
5.9. 1.0.0 Only: NVIDIA GPU Manager for VMware vCenter registration fails with connection refused error.....	16
5.10. 1.0.0 Only: NVIDIA GPU Manager UI fails to open with Plugin Access Check Failed error.....	16
5.11. 1.0.0 Only: NVIDIA GPU Manager for VMware vCenter registration fails with invalid CA certificate error.....	17

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	<ul style="list-style-type: none">▶ 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 <div style="border: 1px solid gray; background-color: #f0f0f0; padding: 5px; margin-top: 10px;"> Note: The base VMware vSphere Hypervisor (ESXi) 7.0 release and 7.0 Update 1 are not supported.</div>
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 1.4

- ▶ Security updates as listed in [Security Updates](#)
- ▶ Miscellaneous bug fixes as listed in [Issues Resolved in Release 1.4](#)

Changes in Release 1.3

- ▶ Support for VMware vSphere 8.0 with Tanzu
- ▶ Security updates as listed in [Security Updates](#)
- ▶ Miscellaneous bug fixes as listed in [Issues Resolved in Release 1.3](#)

Changes in Release 1.2

- ▶ Support for communication through a proxy server
- ▶ Functionality to reset SSL certificates through the web console of the vCenter web client
- ▶ Security updates as listed in [Security Updates](#)
- ▶ Miscellaneous bug fixes as listed in [Issues Resolved in Release 1.2](#)

Changes in Release 1.1

- ▶ Dark theme support for the web-based user interface (UI) to the **NVIDIA GPU Manager** and the **vSphere client plugin**
 - ▶ For **NVIDIA GPU Manager**, you can select the theme from the **NVIDIA GPU Manager** UI.
 - ▶ For the **vSphere client plugin**, the theme matches the theme that is set for the vSphere client.
- ▶ Localization of the **NVIDIA GPU Manager** into the following languages:
 - ▶ Chinese (simplified)
 - ▶ Chinese (traditional)
 - ▶ French

- ▶ German
- ▶ Italian
- ▶ Japanese
- ▶ Korean
- ▶ Spanish
- ▶ Security updates as listed in [Security Updates](#)
- ▶ Miscellaneous bug fixes as listed in [Issues Resolved in Release 1.1](#)

Changes in Release 1.0

Release 1.0 is the first generally available release of NVIDIA GPU Manager for VMware vCenter.

Chapter 3. Security Updates

Security Updates in Release 1.4

To address vulnerabilities CVE-2023-6931, CVE-2023-6932, CVE-2023-51781, and CVE-2024-0646, a new version of the OS package `linux-libc-dev` (5.4.0-172.190) is included in NVIDIA GPU Manager for VMware vCenter.

Security Updates in Release 1.3

- ▶ To address vulnerabilities CVE-2023-4622, CVE-2023-4623, CVE-2023-4921, CVE-2023-42752, CVE-2023-42753, CVE-2023-42755, and CVE-2023-42756, a new version of the OS package `linux-libc-dev` (5.4.0-169.187) is included in NVIDIA GPU Manager for VMware vCenter.
- ▶ To address vulnerability CVE-2023-23583, a new version of the OS package `intel-microcode` (3.20231114.0ubuntu0.20.04.1) is included in NVIDIA GPU Manager for VMware vCenter.
- ▶ To address vulnerability CVE-2023-33733, new versions of the OS packages `python3-renderpm`, `python3-reportlab`, `python3-reportlab-accel` (3.5.34-1ubuntu1.1) are included in NVIDIA GPU Manager for VMware vCenter.

Security Updates in Release 1.2

- ▶ To address vulnerabilities CVE-2023-3609, CVE-2023-3611, CVE-2023-3776, and CVE-2023-20593, a new version of the OS package `linux-libc-dev` (5.4.0-159.176) is included in NVIDIA GPU Manager for VMware vCenter.
- ▶ To address vulnerabilities CVE-2023-20569 and CVE-2023-20593, a new version of the OS package `amd64-microcode` (3.20191218.1ubuntu1.2) is included in NVIDIA GPU Manager for VMware vCenter.

Security Updates in Release 1.1

To address vulnerability CVE-2023-1281, a new version of the OS package `linux-libc-dev` (5.4.0-153.170) is included in NVIDIA GPU Manager for VMware vCenter.

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 1.4

Bug ID	Summary and Description
4507972	<p><u>1.3.0 Only: Static IP address cannot be set during installation of the NVIDIA GPU Manager for VMware vCenter virtual appliance</u></p> <p>A static IP address for an NVIDIA GPU Manager for VMware vCenter virtual appliance cannot be set during installation of the appliance because network settings for a static IP address are missing from the VMware vSphere Client UI.</p>

4.2. Issues Resolved in Release 1.3

Bug ID	Summary and Description
4442922	<p><u>1.0.0-1.2.0 Only: NVIDIA GPU Manager for VMware vCenter registration fails if IPv6 and IPv4 addresses are associated with the vCenter Server domain name</u></p> <p>If an IPv6 address and an IPv4 address are associated with the fully qualified domain name of a vCenter Server host, any attempt to register NVIDIA GPU Manager for VMware vCenter with the host by specifying the host's fully qualified domain name fails. This issue occurs because NVIDIA GPU Manager for VMware vCenter does not support IPv6 addresses and enforces this restriction by checking whether an IPv6 address is associated with the host's fully qualified domain name. If so, registration fails even if an IPv4 address is also associated with the host's fully qualified domain name. When this issue occurs, the following error message is displayed.</p> <p>IPv6 addresses are not supported</p>

4.3. Issues Resolved in Release 1.2

Bug ID	Summary and Description
4117784	<p><u>1.0.0, 1.1.0 Only: NVIDIA GPU Manager does not support communication through a proxy server</u></p> <p>NVIDIA GPU Manager does not support communication through a proxy server. As a result, if a proxy server is used in communications with NVIDIA GPU Manager, the following operations on NVIDIA GPU Manager fail:</p> <ul style="list-style-type: none"> ▶ Registration of NVIDIA GPU Manager fails because VMware vCenter Server cannot reach the NVIDIA GPU Manager drivers repository on the NVIDIA GPU Manager for VMware vCenter virtual appliance. ▶ Connection to the NVIDIA driver download service and API service fails because NVIDIA GPU Manager cannot reach the NVIDIA Licensing Portal.
4334769	<p><u>1.1.0 Only: NVIDIA GPU Manager for VMware vCenter registration fails with incompatible error on vSphere 7.0.3.01700</u></p> <p>As a result of an issue with vCenter Server, vSphere Client cannot find a manifest file that is compatible with its version, namely, 7.0.3.1700. When this issue occurs, vSphere Client shows that deployment of NVIDIA GPU Manager has failed with following error:</p> <pre>The plug-in is incompatible with the current vSphere Client version. com.vmware.vise.plugin.capability.IncompatibleVsphereClientException: Invalid vSphere Client environment. Expected version: '[8.0,)', Actual: '7.0.3.1700'.</pre>

4.4. Issues Resolved in Release 1.1

Bug ID	Summary and Description
4116054	<p><u>1.0.0 Only: NVIDIA GPU Manager UI fails to open with Plugin Access Check Failed error</u></p> <p>The NVIDIA GPU Manager UI fails to open on the vSphere client, which displays the <code>Plugin Access Check Failed</code> error. This issue occurs because the vSphere client is being accessed through the fully qualified domain name of the VMware vCenter Server instance, but the fully qualified domain name is not all in lowercase.</p>
4117722	<p><u>1.0.0 Only: NVIDIA GPU Manager for VMware vCenter registration fails with connection refused error</u></p> <p>If an NVIDIA GPU Manager for VMware vCenter virtual appliance is assigned an IP address that is not mapped to a valid fully qualified domain name, NVIDIA GPU Manager for VMware vCenter registration fails. This error occurs because, by default, NVIDIA GPU Manager uses the fully qualified domain name for communications. If the fully qualified domain</p>

Bug ID	Summary and Description
	name is invalid, the NVIDIA GPU Manager cannot be located through this name, and Error No. 111 Connection refused occurs when registration fails.
4140398	<p data-bbox="500 359 1386 422"><u>1.0.0 Only: NVIDIA GPU Manager for VMware vCenter registration fails with invalid CA certificate error</u></p> <p data-bbox="500 438 1421 564">The SSL certificate that is installed on an NVIDIA GPU Manager for VMware vCenter virtual appliance might not be from a trusted certificate authority (CA) and contain a certificate chain. If you try to register the NVIDIA GPU Manager for VMware vCenter, the registration fails with the following error:</p> <p data-bbox="500 569 1421 623">not a valid CA certificate. Please retry with a valid certificate chain</p>

Chapter 5. Known Issues

5.1. 1.3.0 Only: Static IP address cannot be set during installation of the NVIDIA GPU Manager for VMware vCenter virtual appliance

Description

A static IP address for an NVIDIA GPU Manager for VMware vCenter virtual appliance cannot be set during installation of the appliance because network settings for a static IP address are missing from the **VMware vSphere Client** UI.

Workaround

If the NVIDIA GPU Manager for VMware vCenter virtual appliance can obtain network settings from a DHCP server, complete the installation of the appliance and assign a static IP address to the appliance from the **NVIDIA GPU Manager** application on the appliance or from the hypervisor. For instructions, refer to the following topics in *GPU Manager for VMware vCenter User Guide*:

- ▶ [Updating the Network Settings for the NVIDIA GPU Manager for VMware vCenter Virtual Appliance](#)
- ▶ [Setting the IP Address of an NVIDIA GPU Manager for VMware vCenter Virtual Appliance from the Hypervisor](#)

If only static IP addresses are available in your network environment, revert to the previous release of the NVIDIA GPU Manager for VMware vCenter.

Status

Resolved in NVIDIA GPU Manager for VMware vCenter 1.4.0

Ref.

4507972

5.2. 1.0.0-1.2.0 Only: NVIDIA GPU Manager for VMware vCenter registration fails if IPv6 and IPv4 addresses are associated with the vCenter Server domain name

Description

If an IPv6 address and an IPv4 address are associated with the fully qualified domain name of a vCenter Server host, any attempt to register NVIDIA GPU Manager for VMware vCenter with the host by specifying the host's fully qualified domain name fails. This issue occurs because NVIDIA GPU Manager for VMware vCenter does not support IPv6 addresses and enforces this restriction by checking whether an IPv6 address is associated with the host's fully qualified domain name. If so, registration fails even if an IPv4 address is also associated with the host's fully qualified domain name. When this issue occurs, the following error message is displayed.

```
IPv6 addresses are not supported
```

Workaround

When registering NVIDIA GPU Manager for VMware vCenter, specify the IPv4 address of the vCenter Server host, **not** its fully qualified domain name.

Status

Resolved in NVIDIA GPU Manager for VMware vCenter 1.3.0



Note: Resolution of this issue affects only fully qualified domain names with which both an IPv6 address and an IPv4 address are associated. Because NVIDIA GPU Manager for VMware vCenter does **not** support IPv6 addresses, registration still fails if only an IPv6 address is associated with the fully qualified domain name of the vCenter Server host.

Ref.

4442922

5.3. 1.1.0 Only: NVIDIA GPU Manager for VMware vCenter registration fails with incompatible error on vSphere 7.0.3.01700

Description

As a result of an issue with vCenter Server, vSphere Client cannot find a manifest file that is compatible with its version, namely, 7.0.3.1700. When this issue occurs, vSphere Client shows that deployment of NVIDIA GPU Manager has failed with following error:

```
The plug-in is incompatible with the current vSphere Client version.
com.vmware.vise.plugin.capability.IncompatibleVsphereClientException:
Invalid vSphere Client environment. Expected version: '[8.0,)', Actual:
'7.0.3.1700'.
```

Workaround

1. Use secure shell (SSH) to log in to vCenter Server.
2. In a plain text editor, edit the file `/etc/vmware/vsphere-ui/vc-packages/vsphere-client-serenity/com.nvidia.gpu.manager-1.1.0.0/plugin-80.json` as follows.
 - a). Append a comma (,) to line 16 after the closing brace (}).
 - b). On the next line, add `"open.modal.api": "1.1.0"` with the correct indentation.

```
    },
    "open.modal.api": "1.1.0"
```

3. Save your changes and quit the editor.
4. Restart vSphere Client.

```
service-control --restart vsphere-ui
```



Note: If NVIDIA GPU Manager for VMware vCenter is re-registered, updates to the `plugin-80.json` are lost must be redone.

Status

Resolved in NVIDIA GPU Manager for VMware vCenter 1.2.0

Ref.

4334769

5.4. 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

5.5. 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

5.6. 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

5.7. Since 1.2.0: System Text on the Proxy Server Settings page is not localized

Description

Most of the system text on the **Proxy Server Settings** page, which was added in NVIDIA GPU Manager for VMware vCenter release 1.2.0, is not localized.

Status

Open

Ref.

4326144

5.8. 1.0.0, 1.1.0 Only: NVIDIA GPU Manager does not support communication through a proxy server

Description

NVIDIA GPU Manager does not support communication through a proxy server. As a result, if a proxy server is used in communications with **NVIDIA GPU Manager**, the following operations on **NVIDIA GPU Manager** fail:

- ▶ Registration of **NVIDIA GPU Manager** fails because VMware vCenter Server cannot reach the **NVIDIA GPU Manager drivers repository** on the NVIDIA GPU Manager for VMware vCenter virtual appliance.

- ▶ Connection to the NVIDIA driver download service and API service fails because **NVIDIA GPU Manager** cannot reach the NVIDIA Licensing Portal.

Status

Resolved in NVIDIA GPU Manager for VMware vCenter 1.2.0

Ref.

4117784

5.9. 1.0.0 Only: NVIDIA GPU Manager for VMware vCenter registration fails with connection refused error

Description

If an NVIDIA GPU Manager for VMware vCenter virtual appliance is assigned an IP address that is not mapped to a valid fully qualified domain name, NVIDIA GPU Manager for VMware vCenter registration fails. This error occurs because, by default, **NVIDIA GPU Manager** uses the fully qualified domain name for communications. If the fully qualified domain name is invalid, the **NVIDIA GPU Manager** cannot be located through this name, and `Error No. 111 Connection refused` occurs when registration fails.

Status

Resolved in NVIDIA GPU Manager for VMware vCenter 1.1.0

Ref.

4117722

5.10. 1.0.0 Only: NVIDIA GPU Manager UI fails to open with Plugin Access Check Failed error

Description

The **NVIDIA GPU Manager** UI fails to open on the vSphere client, which displays the `Plugin Access Check Failed` error. This issue occurs because the vSphere client is

being accessed through the fully qualified domain name of the VMware vCenter Server instance, but the fully qualified domain name is not all in lowercase.

Status

Resolved in NVIDIA GPU Manager for VMware vCenter 1.1.0

Ref.

4116054

5.11. 1.0.0 Only: NVIDIA GPU Manager for VMware vCenter registration fails with invalid CA certificate error

Description

The SSL certificate that is installed on an NVIDIA GPU Manager for VMware vCenter virtual appliance might not be from a trusted certificate authority (CA) and contain a certificate chain. If you try to register the NVIDIA GPU Manager for VMware vCenter, the registration fails with the following error:

```
not a valid CA certificate. Please retry with a valid certificate chain
```

Status

Resolved in NVIDIA GPU Manager for VMware vCenter 1.1.0

Ref.

4140398

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, 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

© 2023-2024 NVIDIA Corporation. All rights reserved.

