



# Virtual GPU Management Pack for VMware vRealize Operations

Release Notes


# Table of Contents

- Chapter 1. Supported Software Releases..... 1
- Chapter 2. Changes in this Release..... 2
- Chapter 3. Resolved Issues..... 3
- Chapter 4. Known Issues..... 5
  - 4.1. 2.0, 2.1 Only: The Relationship Tree shows only one vGPU when multiple vGPUs are assigned to a VM..... 5
  - 4.2. 2.0 Only: NVIDIA vGPU adapter stops working after VMware software upgrade..... 6
  - 4.3. 2.0 Only: NVIDIA vGPU adapter fails to identify the vCenter adapter..... 6
  - 4.4. The Search for a vGPU widget lists only one vGPU after navigation from the GPU Summary dashboard..... 7
  - 4.5. 2.0 Only: The GPU Properties widget does not show any data..... 8
  - 4.6. 2.0 Only: NVIDIA vGPU adapter is marked as failed and no data is collected..... 10
  - 4.7. NVIDIA vGPU adapter instance stops collecting data..... 11
  - 4.8. The Alerts on vGPUs running on the selected Host widget is not updated..... 12
  - 4.9. NVIDIA vGPU data is missing from the VMware vRealize Operations dashboards..... 13
  - 4.10. The NVIDIA Host Summary dashboard shows alerts unrelated to the GPU..... 14
  - 4.11. NVIDIA dashboards are not removed after the adapter is uninstalled..... 14

---

# Chapter 1. Supported Software Releases

NVIDIA Virtual GPU Management Pack for VMware vRealize Operations is supported on specific releases of VMware vRealize Operations Manager, VMware vRealize Operations Cloud, and NVIDIA vGPU software.

Software	Supported Releases
VMware vRealize Operations Manager	8.0 through 8.6.3   <b>Note:</b> NVIDIA Virtual GPU Management Pack for VMware vRealize Operations supports only releases of VMware vRealize Operations Manager that are also supported by VMware.
VMware vRealize Operations Cloud	Current generally available release
NVIDIA vGPU software	All releases in all supported release branches

---

# Chapter 2. Changes in this Release

## Changes in Release 2.2

- ▶ Support for VMware vRealize Operations Cloud
- ▶ Miscellaneous bug fixes

## Changes in Release 2.1

- ▶ Security updates
- ▶ Miscellaneous bug fixes

## Changes in Release 2.0

- ▶ Monitoring of multiple vGPUs assigned to the same VM  
NVIDIA Virtual GPU Management Pack for VMware vRealize Operations can now identify configurations in which multiple vGPUs are assigned to the same VM and monitor the metrics and analytics of these configurations.



**Note:** Monitoring of multiple vGPUs assigned to the same VM requires NVIDIA vGPU software 11.0 or later.

- ▶ Support for vRealize Operations Manager 8.0 and later compatible 8.x updates
- ▶ Withdrawal of support for the following vRealize Operations Manager releases:
  - ▶ 7.0 and later compatible 7.x updates
  - ▶ 6.6 and later compatible 6.x updates
- ▶ Miscellaneous bug fixes

---

## Chapter 3. 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 Virtual GPU Management Pack for VMware vRealize Operations **before the issue was resolved**.

### Issues Resolved in Release 2.2

Bug ID	Summary and Description
3666156	<p><b><u>2.0, 2.1 Only: The Relationship Tree shows only one vGPU when multiple vGPUs are assigned to a VM</u></b></p> <p>When multiple vGPUs are assigned to a VM, the <b>Relationship Tree</b> for the VM shows only one vGPU. However, the NVIDIA vGPU adapter creates a separate vGPU object named <i>vm-name-vgpu-profile-name</i> for each vGPU assigned to the VM. These vGPU objects are listed correctly on the <b>NVIDIA vGPU Summary</b> dashboard. This issue occurs because the NVIDIA vGPU adapter sets the relationship between each vGPU and the VM separately. As a result, the previously added relationship between a vGPU and the VM is cleared.</p>

### Issues Resolved in Release 2.1

Bug ID	Summary and Description
3662340	<p><b><u>2.0 Only: NVIDIA vGPU adapter stops working after VMware software upgrade</u></b></p> <p>After an upgrade to vRealize Operations Manager or VMware vSphere ESXi on a host on which NVIDIA Virtual GPU Management Pack for VMware vRealize Operations is installed, the NVIDIA vGPU adapter stops working.</p>
3150425	<p><b><u>2.0 Only: NVIDIA vGPU adapter fails to identify the vCenter adapter</u></b></p> <p>When the software-defined data center (SDDC) Health Monitoring adapter is installed on a vRealize Operations node, the NVIDIA vGPU adapter sometimes fails to identify the vCenter adapter instance. This issue occurs because the NVIDIA vGPU adapter needs the vCenter adapter to be in the data receiving state to run.</p>

Bug ID	Summary and Description
3050041	<p><b><u>2.0 Only: NVIDIA vGPU adapter is marked as failed and no data is collected</u></b></p> <p>The NVIDIA vGPU adapter is marked as failed and no data is collected. When this issue occurs, the error <code>isVCenterAdapterInstanceRunning - Error: ResourceDto object returned by vROPs API is null</code> is reported by VMware vCenter Server.</p>
200702479	<p><b><u>2.0 Only: The GPU Properties widget does not show any data</u></b></p> <p>The <b>GPU Properties</b> widget in the <b>NVIDIA GPU Summary</b> dashboard does not show any data. This issue occurs because the case of an identifier in the <b>GPU Utilization is high</b> alert definition is incorrect. Starting with vRealize Operations Manager release 8.3, identifiers are case-sensitive.</p>

### Issues Resolved in Release 2.0

Bug ID	Summary and Description
200525085	<p><b>NVIDIA vGPU adapter instance does not collect data after VM restart</b></p> <p>After a VMware vRealize Operations VM is restarted, the NVIDIA vGPU adapter instance does not start collecting data.</p>
200635615	<p><b>NVIDIA vGPU adapter objects flood the vRealize Operations Manager instance</b></p> <p>NVIDIA vGPU adapter objects of type <b>Process</b> can flood the vRealize Operations Manager instance. When the number of objects becomes too high for the vRealize Operations Manager instance to process, performance degradation and usability problems might be observed.</p>

---

## Chapter 4. Known Issues

### 4.1. 2.0, 2.1 Only: The Relationship Tree shows only one vGPU when multiple vGPUs are assigned to a VM

#### Description

When multiple vGPUs are assigned to a VM, the **Relationship Tree** for the VM shows only one vGPU. However, the NVIDIA vGPU adapter creates a separate vGPU object named *vm-name-vgpu-profile-name* for each vGPU assigned to the VM. These vGPU objects are listed correctly on the **NVIDIA vGPU Summary** dashboard. This issue occurs because the NVIDIA vGPU adapter sets the relationship between each vGPU and the VM separately. As a result, the previously added relationship between a vGPU and the VM is cleared.

#### Status

Resolved in NVIDIA Virtual GPU Management Pack for VMware vRealize Operations 2.2

#### Ref. #

3666156

## 4.2. 2.0 Only: NVIDIA vGPU adapter stops working after VMware software upgrade

### Description

After an upgrade to vRealize Operations Manager or VMware vSphere ESXi on a host on which NVIDIA Virtual GPU Management Pack for VMware vRealize Operations is installed, the NVIDIA vGPU adapter stops working.

### Workaround

Uninstall and reinstall NVIDIA Virtual GPU Management Pack for VMware vRealize Operations.

### Status

Resolved in NVIDIA Virtual GPU Management Pack for VMware vRealize Operations 2.1

### Ref. #

3662340

## 4.3. 2.0 Only: NVIDIA vGPU adapter fails to identify the vCenter adapter

### Description

When the software-defined data center (SDDC) Health Monitoring adapter is installed on a vRealize Operations node, the NVIDIA vGPU adapter sometimes fails to identify the vCenter adapter instance. This issue occurs because the NVIDIA vGPU adapter needs the vCenter adapter to be in the data receiving state to run.

When this issue occurs, error messages similar to the following examples are logged.

```
2020-10-01T15:38:35,706 ERROR [Task Processor worker thread 5] (4973936)
com.integrien.alive.common.adapter3.AdapterBase.applyConfiguration
- Failed to configure adapter 'NvVGPUAdapter': VMWARE vCenter adapter is either
not
configured or not receiving data for VC: server.example.com
2020-10-01T15:55:49,853 DEBUG [Task Processor worker thread 2] (4973936)
com.nvidia.nvvgpu.adapter.client.VropsInterface.isVCenterAdapterInstanceRunning
- Error: vCenter adapter collection status is NONE. Re-checking in 2 mins
[Repeated 8 times ...]
```



## Status

Resolved in NVIDIA Virtual GPU Management Pack for VMware vRealize Operations 2.1

## Ref. #

3150425

# 4.4. The Search for a vGPU widget lists only one vGPU after navigation from the GPU Summary dashboard

## Description

After a user navigates from the **GPU Summary** dashboard to the **vGPU Summary** dashboard, the **Search for a vGPU** widget lists only one vGPU. This issue occurs when the user navigates between the dashboards by using the navigation button in the **vGPUs running in selected GPU** widget. When this issue occurs, the **Search for a vGPU** widget lists only the vGPU that was selected in the **vGPUs running in selected GPU** widget.

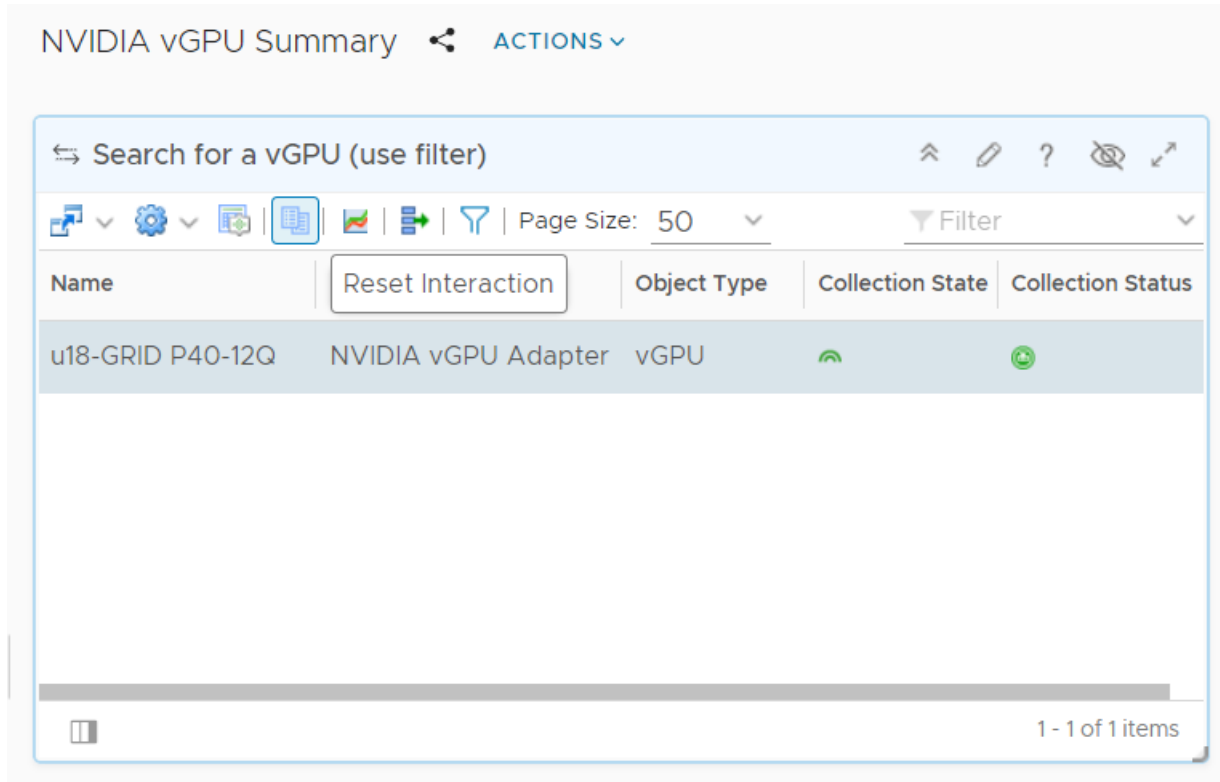
This issue occurs because the concept of dashboard-to-dashboard navigation was changed in vRealize Operations Manager release 8.3.

## Version

This issue affects vRealize Operations Manager release 8.3 and later 8.x updates.

## Workaround

In the **Search for a vGPU** widget on the **vGPU Summary** dashboard, click **Reset Interaction**.



All the vGPUs present are now listed.

### Status

Not an NVIDIA bug

### Ref. #

200702483

## 4.5. 2.0 Only: The GPU Properties widget does not show any data

### Description

The **GPU Properties** widget in the **NVIDIA GPU Summary** dashboard does not show any data. This issue occurs because the case of an identifier in the **GPU Utilization is high** alert definition is incorrect. Starting with vRealize Operations Manager release 8.3, identifiers are case-sensitive.

## Version

This issue affects vRealize Operations Manager release 8.3 and later 8.x updates.

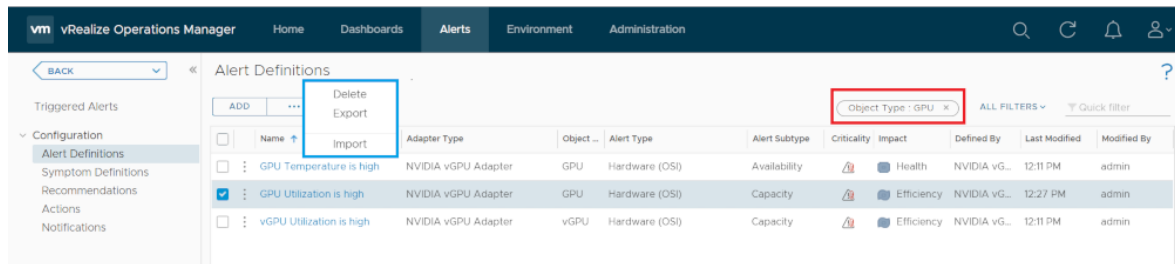
## Workaround

Use the supplied alert definition file to update the alert definition to use the correct case.

1. Download the [gpu-utilization-is-high.zip](#) file and extract the `gpu-utilization-is-high.xml` file that it contains.

Ensure that the extracted file is accessible to the web browser that you are using to manage your **vRealize Operations Manager** instance.

2. Log in to your **vRealize Operations Manager** instance as an administrator user.
3. On the **vRealize Operations Manager Home** page, follow the **Alerts** link.
4. In the navigation bar, expand **Configuration** and select **Alert Definitions**.
5. On the **Alert Definitions** page that opens, delete the **GPU Utilization is high** alert.
  - a). Set the **Object Type: GPU** filter.
  - b). Select the **GPU Utilization is high** alert and click **Delete**.



6. Import a new **GPU Utilization is high** alert definition from the file `gpu-utilization-is-high.xml`.
  - a). On the **Alert Definitions** page, click **Import**.
  - b). In the **Import Alert Definition** window that opens, set the **Overwrite existing Alert Definition** option.
  - c). Click **BROWSE** and navigate to and select the file `gpu-utilization-is-high.xml`.

## Import Alert Definition

Select an Alert Definition XML file to import.

**BROWSE...**

**i** The import process begins when you select the file in the browse window.

In case of a conflict:

Overwrite existing Alert Definition

Skip import

**CANCEL**

7. When the import process is complete, click **DONE**.

It might take up to 15 minutes after the import process is complete for the dashboard to be updated.

### Status

Resolved in NVIDIA Virtual GPU Management Pack for VMware vRealize Operations 2.1

### Ref. #

200702479

## 4.6. 2.0 Only: NVIDIA vGPU adapter is marked as failed and no data is collected

### Description

The NVIDIA vGPU adapter is marked as failed and no data is collected. When this issue occurs, the error `isvCenterAdapterInstanceRunning - Error: ResourceDto object returned by vROPs API is null` is reported by VMware vCenter Server.

When this error is reported, the VMware vCenter Server adapter object that is returned is `NULL`. As a result, the NVIDIA vGPU adapter can not proceed and exits. The VMware vCenter Server adapter object is used to check if the VMware vCenter Server adapter is receiving data.

The VMware vCenter Server configuration in the NVIDIA vGPU adapter is case sensitive. This issue occurs if the case of the VMware vCenter Server name in the NVIDIA vGPU adapter and

the VMware vCenter Server adapter do not match. For example, this issue occurs if the NVIDIA vGPU adapter configuration specifies `vcenter.example.com` and the VMware vCenter Server adapter configuration specifies `VCENTER.example.com`.

## Workaround

Change case of the VMware vCenter Server name in the NVIDIA vGPU adapter to match the case in the VMware vCenter Server adapter.

## Status

Resolved in NVIDIA Virtual GPU Management Pack for VMware vRealize Operations 2.1

## Ref. #

3050041

# 4.7. NVIDIA vGPU adapter instance stops collecting data

## Description

After some data collection cycles, the NVIDIA vGPU adapter instance randomly stops collecting data.

When this issue occurs, the following errors are written to the NVIDIA vGPU adapter log file:

```
Collector worker thread 25] (13350)
  com.nvidia.nvvgpu.adapter.client.DcgmClient.getHostConfig - Starting collection for
  host: 10.24.131.52
[30740] 2019-01-18 11:47:45,414 DEBUG [Collector worker thread 25] (13350)
  com.nvidia.nvvgpu.adapter.client.DcgmClient.getGroupInfo - Sending DCGM Command:
  GROUPINFO
[30741] 2019-01-18 11:48:03,805 DEBUG [pool-868-thread-1] (13350)
  com.nvidia.nvvgpu.adapter.client.CimClient.run - Retrieving hosts and initializing
  CIM Client instances
[30742] 2019-01-18 11:48:22,111 ERROR [pool-868-thread-1] (13350)
  com.nvidia.nvvgpu.adapter.client.CimClient.run - java.lang.RuntimeException:
  java.rmi.RemoteException:
  VI SDK invoke exception:java.net.UnknownHostException: dc4dvvc01.nvidia.com
```

An error similar to the following example is also written to the NVIDIA vGPU log files, the `/var/log/messages` file, or the `syslog` file for all the hosts that are reporting failure:

```
Timeout error accepting SSL connection
```

The root cause of this issue is a known issue with VMware vSphere Hypervisor (ESXi). For more information, see [VMware Knowledge Base Article: VMware ESX/ESXi host logs timeout errors when trying to establish SSL connections \(1020806\)](#).

## Workaround

1. In a plain-text editor, open the configuration file for the `sfc` service `/etc/sfc/sfc.cfg` on the host where the adapter stopped collecting data.
2. Change the value of the property `httpsProcs` to 8.
3. Save your changes and quit the editor.
4. Restart the `sfc` service.

## Status

Not an NVIDIA bug

## Ref. #

200486366

# 4.8. The Alerts on vGPUs running on the selected Host widget is not updated

## Description

The **Alerts on vGPUs running on the selected Host** widget on the **NVIDIA Host Summary** dashboard is not updated. This issue affects **only** the **NVIDIA Host Summary** dashboard. The **NVIDIA GPU Summary** dashboard and the **NVIDIA vGPU Summary** dashboard are updated with the relevant alerts.

## Workaround



**Note:** This workaround does **not** work on vRealize Operations Manager 7.5 or later releases.

Edit and save the **Alerts on vGPUs running on the selected Host** widget on the **NVIDIA Host Summary** dashboard.

## Status

Not an NVIDIA bug

## Ref. #

200344549

## 4.9. NVIDIA vGPU data is missing from the VMware vRealize Operations dashboards

### Description

To collect data from hosts in VMware vCenter that are running NVIDIA GPUs and the NVIDIA vGPU Manager, each user of the NVIDIA vGPU adapter requires the CIM interaction privilege. If this privilege is not assigned, the user cannot use the NVIDIA vGPU adapter to collect data.

When this issue occurs, the adapter log files contain error messages similar to the following examples:

```
2019-07-01 17:40:32,296 DEBUG [pool-9771-thread-1] (117)
  com.nvidia.nvvgpu.adapter.client.CimClient.initializeWBEMClient -
  com.vmware.vim25.NoPermission
2019-07-01 17:40:32,296 WARN [pool-9771-thread-1] (117)
  com.nvidia.nvvgpu.adapter.client.CimClient.initializeWBEMClient - CIM Connection to
  host: srvr-12.example.com failed. This host will be skipped from current collection
  cycle
2019-07-01 17:41:32,296 DEBUG [pool-9771-thread-1] (117)
  com.nvidia.nvvgpu.adapter.client.CimClient.run - Retrieving hosts and initializing
  CIM Client instances
2019-07-01 17:41:32,328 INFO [pool-9771-thread-1] (117)
  com.nvidia.nvvgpu.adapter.client.CimClient.initializeWBEMClient - Initializing CIM
  Client for host: srvr-10.example.com
2019-07-01 17:41:32,330 DEBUG [pool-9771-thread-1] (117)
  com.nvidia.nvvgpu.adapter.client.CimClient.initializeWBEMClient -
  com.vmware.vim25.NoPermission
2019-07-01 17:41:32,331 WARN [pool-9771-thread-1] (117)
  com.nvidia.nvvgpu.adapter.client.CimClient.initializeWBEMClient - CIM Connection to
  host: srvr-10.example.com failed. This host will be skipped from current collection
  cycle
2019-07-01 17:41:32,343 INFO [pool-9771-thread-1] (117)
  com.nvidia.nvvgpu.adapter.client.CimClient.initializeWBEMClient - Initializing CIM
  Client for host: srvr-11.example.com
2019-07-01 17:41:32,346 DEBUG [pool-9771-thread-1] (117)
  com.nvidia.nvvgpu.adapter.client.CimClient.initializeWBEMClient -
  com.vmware.vim25.NoPermission
2019-07-01 17:41:32,346 WARN [pool-9771-thread-1] (117)
  com.nvidia.nvvgpu.adapter.client.CimClient.initializeWBEMClient - CIM Connection to
  host: srvr-11.example.com failed. This host will be skipped from current collection
  cycle
2019-07-01 17:41:32,359 INFO [pool-9771-thread-1] (117)
  com.nvidia.nvvgpu.adapter.client.CimClient.initializeWBEMClient - Initializing CIM
  Client for host: srvr-12.example.com
2019-07-01 17:41:32,362 DEBUG [pool-9771-thread-1] (117)
  com.nvidia.nvvgpu.adapter.client.CimClient.initializeWBEMClient -
  com.vmware.vim25.NoPermission
```

### Workaround

Assign the CIM interaction privilege that the NVIDIA vGPU adapter requires.

### Status

Not a bug.

### Ref. #

2639301

## 4.10. The NVIDIA Host Summary dashboard shows alerts unrelated to the GPU

### Description

After the NVIDIA Virtual GPU Management Pack for VMware vRealize Operations is installed, an NVIDIA vGPU adapter instance is created and the host is rebooted, the **NVIDIA Host Summary** dashboard shows alerts unrelated to the GPU.

### Status

Not an NVIDIA bug

### Ref. #

200451772

## 4.11. NVIDIA dashboards are not removed after the adapter is uninstalled

### Description

After the NVIDIA vGPU adapter is uninstalled, NVIDIA dashboards are still present. These dashboards should be removed as a part of the uninstallation process.

### Status

Not an NVIDIA bug

### Ref. #

200343762



## 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

© 2017-2022 NVIDIA Corporation. All rights reserved.