



Virtual GPU Management Pack for VMware Aria Operations



Release Notes

Table of Contents

Chapter 1. Supported Software Releases.....	1
Chapter 2. Changes in this Release.....	2
Chapter 3. Resolved Issues.....	4
Chapter 4. Known Issues.....	5
4.1. GPU Instance Properties widget lists properties for time-sliced vGPUs as ?.....	5
4.2. Compute Instances List widget doesn't list compute instances correctly.....	6
4.3. The nvdGpuMgmtDaemon daemon is killed when multiple VMware Aria Operations instances are collecting data.....	8
4.4. The Search for a vGPU widget lists only one vGPU after navigation from the GPU Summary dashboard.....	8
4.5. NVIDIA vGPU adapter instance stops collecting data.....	10
4.6. The Alerts on vGPUs running on the selected Host widget is not updated.....	11
4.7. NVIDIA vGPU data is missing from the VMware vRealize Operations dashboards.....	11
4.8. The NVIDIA Host Summary dashboard shows alerts unrelated to the GPU.....	13
4.9. NVIDIA dashboards are not removed after the adapter is uninstalled.....	13

Chapter 1. Supported Software Releases

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

Software	Supported Releases
VMware Aria Operations Manager	<p>Since 3.2: 8.0 through 8.12</p> <p>3.1 only: 8.0 through 8.10</p> <div> Note: NVIDIA Virtual GPU Management Pack for VMware Aria Operations supports only releases of VMware Aria Operations Manager that are also supported by VMware.</div>
VMware Aria Operations Cloud	Current generally available release
NVIDIA vGPU software	<p>All releases in all supported release branches except NVIDIA vGPU software 15.0</p> <div> Note: NVIDIA vGPU software 15.0 is not supported because support for Virtual GPU Manager for VMware vSphere versions based on the VMware Daemon SDK (DSDK) is introduced in NVIDIA vGPU software release 15.1.</div>

Chapter 2. Changes in this Release



Note: VMware has changed the product name of its management platform software from VMware vRealize Operations Manager to **VMware Aria Operations Manager**.

Changes in Release 3.2

Metrics for MIG-backed vGPUs are now supported.



Note:

- ▶ Metrics for MIG-backed vGPUs are **not** supported by NVIDIA vGPU software releases before 15.1.
- ▶ Engine utilization metrics are not supported on MIG devices, namely, MIG-enabled GPUs and MIG-backed vGPUs.

Changes in Release 3.1

The issue that caused NVIDIA to withdraw NVIDIA Virtual GPU Management Pack for VMware Aria Operations 3.0 has been resolved.

Changes in Release 3.0



Note: NVIDIA has withdrawn NVIDIA Virtual GPU Management Pack for VMware Aria Operations 3.0 after becoming aware of an issue with this release that causes data collection to fail after an upgrade from release 2.2.

- ▶ Virtual GPU Manager for VMware vSphere versions based on the VMware Daemon SDK (DSDK) are supported starting with NVIDIA vGPU software release 15.1.
- ▶ Required classes for debugging are no longer added to the VMware Aria Operations logs by default.

Instead, you must explicitly add these classes to the VMware Aria Operations logs as explained in [Virtual GPU Management Pack for VMware Aria Operations User Guide](#).

- ▶ Security updates are included.

- Miscellaneous bugs have been fixed.

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 Aria Operations **before the issue was resolved**.

Issues Resolved in Release 3.2

No resolved issues are reported in this release of NVIDIA Virtual GPU Management Pack for VMware Aria Operations.

Issues Resolved in Release 3.1

The issue that caused NVIDIA to withdraw NVIDIA Virtual GPU Management Pack for VMware Aria Operations 3.0 has been resolved.

Issues Resolved in Release 3.0


No resolved issues are reported in this release of NVIDIA Virtual GPU Management Pack for VMware Aria Operations.

Chapter 4. Known Issues

4.1. GPU Instance Properties widget lists properties for time-sliced vGPUs as ?

Description

In VMware Aria Operations Manager releases 8.0 and 8.1, the **GPU Instance Properties** widget lists properties for time-sliced vGPUs as a ? character. For time-sliced vGPUs, the **GPU Instance Properties** widget should be empty because GPU instances are specific to MIG-backed vGPUs.



Object Name	Property Name	Value
Win11SV2-1-GRID ...	GPU Instance Id	?
Win11SV2-1-GRID ...	GPU Instance Name	?
Win11SV2-1-GRID ...	GPU Instance Profile Id	?
Win11SV2-1-GRID ...	GPU Instance Slice Count	?
Win11SV2-1-GRID ...	SM Count	?
Win11SV2-1-GRID ...	Copy Engine Count	?

1 - 10 of 12 items

Version

This issue affects VMware Aria Operations Manager releases 8.0 and 8.1.

Workaround

Ignore the ? character that is displayed. In VMware Aria Operations Manager releases 8.0 and 8.1, absent metrics are shown as a ? character. This behavior does not affect the functionality of VMware Aria Operations.

Status

Not an NVIDIA bug

Resolved by VMware in VMware Aria Operations Manager release 8.2.

4.2. Compute Instances List widget doesn't list compute instances correctly

Description

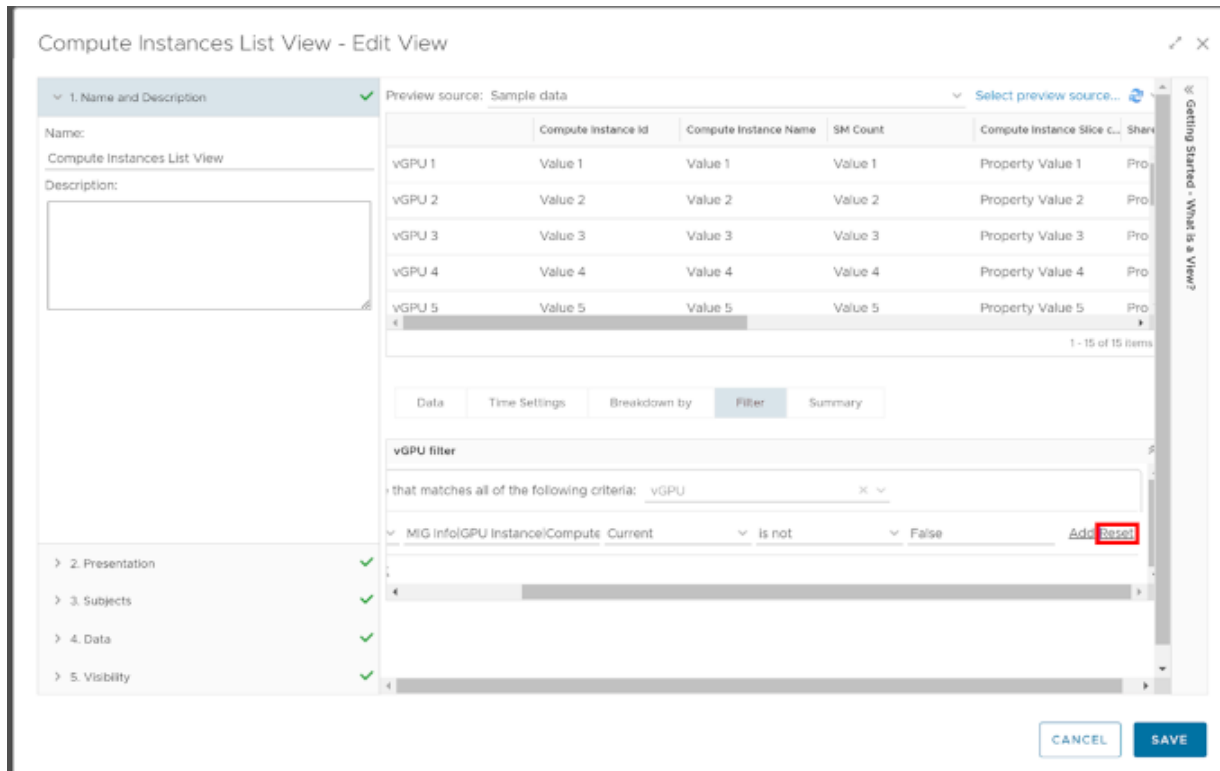
In VMware Aria Operations Manager releases 8.0 and 8.1, the **Compute Instances List** widget doesn't list compute instances correctly. This issue occurs because the **Compute Instances List** widget depends on a feature that was added to VMware Aria Operations Manager 8.2 for filtering instanced metrics and properties of active compute instances. Because this feature is not available in VMware Aria Operations Manager releases 8.0 and 8.1, the **Compute Instances List** widget in these releases cannot list compute instances correctly.

Version

This issue affects VMware Aria Operations Manager releases 8.0 and 8.1.

Workaround

Clear the vGPU filter in the **Compute Instances List** widget.



1. At the top right corner of the **Compute Instances List View** page, click **Edit Widget**.
2. Navigate to **Output Data > Compute Instance List View > Edit**.
3. On the **Compute Instances List View**, follow the **Reset** under the vGPU filter and click **SAVE**.

After the vGPU filter is cleared, the **Compute Instances List View** page listing all active and inactive compute instances. To differentiate between active and inactive compute instances, use the **Compute Instance Alive** option.

Status

Not an NVIDIA bug

Resolved by VMware in VMware Aria Operations Manager release 8.2.

4.3. The `nvdGpuMgmtDaemon` daemon is killed when multiple VMware Aria Operations instances are collecting data

Description

The `nvdGpuMgmtDaemon` daemon is killed when multiple VMware Aria Operations instances are collecting data from a single NVIDIA vGPU host. This issue does not occur when only one VMware Aria Operations instance is collecting data from the NVIDIA vGPU host. When the daemon is killed, GPU data collection fails.

Workaround

Restart the `nvdGpuMgmtDaemon` manually from the ESXi host to resume data collection.

Status

Open

Ref.

4037267

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.

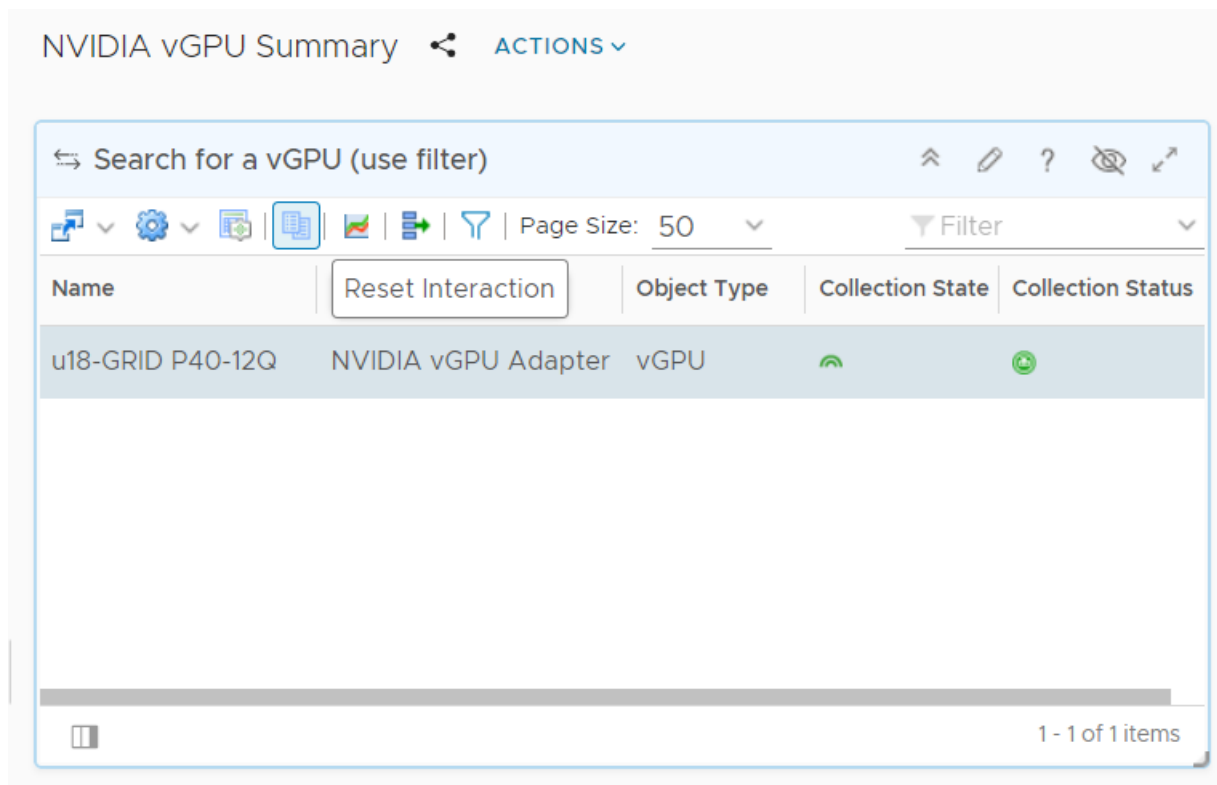
This issue occurs because the concept of dashboard-to-dashboard navigation was changed in vRealize Operations Manager release 8.3.

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. 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

4.6. 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

4.7. 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 an NVIDIA vGPU Manager version that uses a CIM provider, 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

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

Description

After the NVIDIA Virtual GPU Management Pack for VMware Aria 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

4.9. 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

© 2017-2024 NVIDIA Corporation. All rights reserved.

