



VIRTUAL GPU MANAGEMENT PACK FOR VMWARE VREALIZE OPERATIONS

RN-09267-001 _v1.0 through 1.1 | May 2020

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.....	4
4.1. NVIDIA vGPU adapter instance stops collecting data.....	4
4.2. The Alerts on vGPUs running on the selected Host widget is not updated.....	5
4.3. NVIDIA vGPU data is missing from the VMware vRealize Operations dashboards.....	5
4.4. NVIDIA vGPU adapter instance does not collect data after VM restart.....	7
4.5. The NVIDIA Host Summary dashboard shows alerts unrelated to the GPU.....	7
4.6. NVIDIA dashboards are not removed after the adapter is uninstalled.....	8

Chapter 1.

SUPPORTED SOFTWARE RELEASES

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

Software	Supported Releases
vRealize Operations Manager	7.0 and later compatible 7.x updates 6.6 and later compatible 6.x updates
NVIDIA vGPU software	5.0 or later

Chapter 2.

CHANGES IN THIS RELEASE

Changes in Release 1.1

- ▶ New frame buffer statistics for GPU and vGPU objects
- ▶ New **Framebuffer Statistics** widget in the **GPU Summary** and **vGPU Summary** dashboards for displaying the new frame buffer statistics
- ▶ 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 1.1

Bug ID	Summary and Description
200422318	<p>No data displayed on NVIDIA dashboards in VMware vRealize Operations</p> <p>Data is not collected for any ESXi host that is identified by its host name when it is added to VMware vCenter Server. The connection to CIM provider from NVIDIA Virtual GPU Management Pack fails for such hosts because their host names cannot be resolved.</p>
200402630	<p>FrameBuffer Usage chart in the NVIDIA vGPU Summary dashboard always shows the value as zero</p> <p>The FrameBuffer Usage metric in the NVIDIA vGPU Summary dashboard always shows the value as zero irrespective of the actual usage as reported by the <code>nvidia-smi vgpu -q</code> command.</p>
2109975	<p>Connection to vCenter server fails with InvalidLogin because of multiple active sessions</p> <p>When a user tries to configure the NVIDIA vGPU adapter, the following error frequently occurs: Unable to establish a valid connection to the target system. null. Sometimes the connection succeeds, but when the user tries to save configuration settings, the following error occurs: Unable to establish a valid connection to the target system.</p>

Chapter 4.

KNOWN ISSUES

4.1. 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
```

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

Open

Ref. #

200486366

4.2. 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

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

4.3. 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

4.4. NVIDIA vGPU adapter instance does not collect data after VM restart

Description

After a VMware vRealize Operations VM is restarted, the NVIDIA vGPU adapter instance does not start collecting data.

Workaround

Wait until the VMware vCenter adapter shows the collection status as **Data receiving** and restart the NVIDIA vGPU adapter instance.

Status

Open

Ref.

200525085

4.5. 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

4.6. 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 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

© 2017-2020 NVIDIA Corporation. All rights reserved.