



DATACENTER GPU MANAGER 1.6.3

v1.6.3 | April 2019

Release Notes



TABLE OF CONTENTS

Changelog..... iii

 New Features..... iii

 Improvements..... iii

 Bug Fixes..... iv

 Known Issues..... iv

CHANGELOG

New Features

This version of DCGM (v1.6.3) requires a minimum R384 driver that can be downloaded from [NVIDIA Drivers](#). On NVSwitch based systems such as DGX-2 or HGX-2, a minimum of R410 driver is required. It is recommended to install the latest Tesla driver from NVIDIA drivers for use with DCGM.

General

- ▶ DCGM Diagnostics now include a training mode (**--train**). This runs several iterations of the diagnostic to produce a configure file with tuned pass/fail criteria for stable systems. This feature is beta for v1.6.
- ▶ Added a command line parameter (**--throttle-mask**) and a global config file parameter (throttle-mask) for specifying throttle errors to ignore when running a DCGM Diagnostic (**dcgmi diag**)

Platform Support

- ▶ Added DCGM Diagnostics (NVVS) support for Tesla T4, including CUDA 10 compatible plugins.

Improvements

General

- ▶ DCGM Health Watch now performs additional GPU memory health checks.
- ▶ Updated error messages to provide more information about what went wrong when a failure occurs during a DCGM Diagnostic.
- ▶ NVLink bandwidth now reports is MB/s instead of total number of sent bytes since the monitoring start.
- ▶ DCGM can now detach and reattach to GPUs without being restarted.
- ▶ DCGM ignores throttling failures by default on Tesla K80 and T4 GPUs.

Bug Fixes

- ▶ Fixed an issue with debug options in DCGM Diagnostics (e.g. **-debugLogFile**)
- ▶ Fixed an issue where running nv-hostengine without root privileges would return an error when creating a pid file and may hang in some cases.
- ▶ Fixed an issue where DCGM Diagnostics would fail all software tests when the first test failed.
- ▶ **dcgmi diag -h** has been updated to be more descriptive for the run (**-r**) option.
- ▶ Fixed an issue where sending **SIGINT** (e.g. **ctrl-c**) to DCGM diagnostics (**dcgmi diag**) would result in a 0 exit status. Sending a **SIGINT** will now result in a **UNIX SIGINT** signal exit status (130).
- ▶ **dcgmi diag** level one test failures include all failure details from DCGM diagnostics (NVVS).
- ▶ Health watches fixed to now report a full list of failures if multiple failures are present at once.
- ▶ DCGM Diagnostic will only fail on CRC errors if the user configures it (1.5.* would fail with any CRC errors during a test).
- ▶ Fixed an issue with **DcgmHealthWatch::SetResponseV2** API that would in some cases return incorrect values for watch incidents.
- ▶ DCGM now returns the NVML version is returned instead of the driver version.
- ▶ **dcgmi topo** will not report an error if run with a single GPU in the system.
- ▶ Fixed long timeout during a connection to **nv-hostengine** daemon.
- ▶ Fixed a crash in DCGM if a module could not be loaded successfully.
- ▶ Fixed NVLink bandwidth reporting issues (incorrect units in some cases) in DCGM. Note that this fix needs a minimum driver version of 418.40.03.
- ▶ Fixed an issue where the targeted power test (DCGM diagnostics) would incorrectly fail in some cases on Tesla T4 GPUs.

Known Issues

- ▶ DCGM does not support systems with more than 16 GPUs. Support will be added in a future release.
- ▶ With DCGM Diagnostics (**dcgmi diag --help**), the debug log file generation option (**--debugLogFiledebug**) is incorrectly specified in the help listing. The correct option is **--debugLogFile**
- ▶ On P4 based systems, the memory bandwidth test may fail due to not meeting the expected threshold. This is a known issue and to reduce the bandwidth expectations and allow DCGM to succeed, use the following command: **dcgmi diag -r "memory bandwidth" -p "memory bandwidth.minimum_bandwidth=15000"**
- ▶ On K80s, **nvidia-smi** may report hardware throttling (**clocks_throttle_reasons.hw_slowdown = ACTIVE**) during DCGM Diagnostics (Level 3). The stressful workload results in power transients that

engage the HW slowdown mechanism to ensure that the Tesla K80 product operates within the power capping limit for both long term and short term timescales. For Volta or later Tesla products, this reporting issue has been fixed and the workload transients are no longer flagged as "HW Slowdown". The NVIDIA driver will accurately detect if the slowdown event is due to thermal thresholds being exceeded or external power brake event. It is recommended that customers ignore this failure mode on Tesla K80 if the GPU temperature is within specification.

Notice

THE INFORMATION IN THIS GUIDE AND ALL OTHER INFORMATION CONTAINED IN NVIDIA DOCUMENTATION REFERENCED IN THIS GUIDE IS PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE INFORMATION FOR THE PRODUCT, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the product described in this guide shall be limited in accordance with the NVIDIA terms and conditions of sale for the product.

THE NVIDIA PRODUCT DESCRIBED IN THIS GUIDE IS NOT FAULT TOLERANT AND IS NOT DESIGNED, MANUFACTURED OR INTENDED FOR USE IN CONNECTION WITH THE DESIGN, CONSTRUCTION, MAINTENANCE, AND/OR OPERATION OF ANY SYSTEM WHERE THE USE OR A FAILURE OF SUCH SYSTEM COULD RESULT IN A SITUATION THAT THREATENS THE SAFETY OF HUMAN LIFE OR SEVERE PHYSICAL HARM OR PROPERTY DAMAGE (INCLUDING, FOR EXAMPLE, USE IN CONNECTION WITH ANY NUCLEAR, AVIONICS, LIFE SUPPORT OR OTHER LIFE CRITICAL APPLICATION). NVIDIA EXPRESSLY DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY OF FITNESS FOR SUCH HIGH RISK USES. NVIDIA SHALL NOT BE LIABLE TO CUSTOMER OR ANY THIRD PARTY, IN WHOLE OR IN PART, FOR ANY CLAIMS OR DAMAGES ARISING FROM SUCH HIGH RISK USES.

NVIDIA makes no representation or warranty that the product described in this guide will be suitable for any specified use without further testing or modification. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to ensure the product is suitable and fit for the application planned by customer and to do 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 guide. NVIDIA does not accept any 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 guide, or (ii) customer product designs.

Other than the right for customer to use the information in this guide with the product, no other license, either expressed or implied, is hereby granted by NVIDIA under this guide. Reproduction of information in this guide is permissible only if reproduction is approved by NVIDIA in writing, is reproduced without alteration, and is accompanied by all associated conditions, limitations, and notices.

Trademarks

NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the United States and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2013-2019 NVIDIA Corporation. All rights reserved.