



DATACENTER GPU MANAGER 1.7

v1.7 | September 2019

Release Notes



TABLE OF CONTENTS

- Changelog..... iii
- Patch Releases..... iii
 - DCGM v1.7.2..... iii
- DCGM v1.7 GA..... iv
 - New Features..... iv
 - Improvements..... iv
 - Bug Fixes..... iv
 - Known Issues..... v

CHANGELOG

This version of [DCGM](#) (v1.7) 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 R418 driver is required. If using the new profiling metrics capabilities in DCGM, then a minimum of R418 driver is required. It is recommended to install the latest Tesla driver from NVIDIA drivers for use with DCGM.

Patch Releases

DCGM v1.7.2

DCGM v1.7.2 released in December 2019.

Improvements

- ▶ Added support for Quadro RTX 8000 and Quadro RTX 6000.
- ▶ Added support for Tesla V100S-PCIE-32GB.
- ▶ Make the passive health watches (controlled by `dcgmi health`) warn for pending page retirements. They used to report a failure, but warn instead as this failure doesn't prevent the workload from executing.
- ▶ Added the ability to pause and resume DCGM profiling metrics so that profiling can be done while monitoring is enabled. This is done via `dcgmi profile --pause/--resume`.
- ▶ Enabled the NVLink Rx+Tx profiling fields (1011-1012).
- ▶ Added the `dcgmi dmon --nowatch` option to allow `dcgmi` to observe metrics that were already watched by other DCGM clients without affecting the watch frequency or quota policy.

Bug Fixes

- ▶ Fixed the DCGM profiling data sometimes appearing under the wrong GPU in pass-through mode. This could occur if the PCI BDF of the GPUs was changed as the GPUs were passed through.
- ▶ Fixed the first value returned always being 0 for DCP fields 1001-1012. DCP now records a valid value immediately after the fields are watched.
- ▶ Fixed DCP PCIe bandwidth being off by a factor of 2-5x when metrics were multiplexed.

DCGM v1.7 GA

DCGM v1.7.1 released in September 2019.

New Features

General

- ▶ DCGM now supports new profiling metrics at the device-level from GPUs that can be used to understand application behavior. This capability is supported as beta on Linux x86_64 and POWER (ppc64le) platforms. See the User Guide for more information. Note that automatic multiplexing of metrics is alpha.
- ▶ Samples and bindings have been moved to `/usr/local/dcgm`.

Improvements

General

- ▶ DCGM 1.7 requires a minimum `glibc` version of 2.14. As a result, the installation of DCGM on older Linux distributions such as Red Hat Enterprise Linux (RHEL) 6.x or CentOS 6.x may result in an error. See the Supported Platforms section in the User Guide for the minimum system requirements.
- ▶ Added error codes and messages for various DCGM health checks
- ▶ Added a new CLI option `fail-early` to DCGM Diagnostics. This option enables early failure checks for the Targeted Power, Targeted Stress, SM Stress, and Diagnostic tests to check for a failure while the test is running instead at the end of the tests, providing feedback on GPU state quicker to the user
- ▶ Updated error reporting to indicate failures in the CUDA tests when running the MemoryBandwidth tests
- ▶ DCGM documentation can now be found online at <http://docs.nvidia.com/datacenter/dcgm> and packages no longer include documentation.

Bug Fixes

- ▶ The Memory Bandwidth test threshold for P4 products has been changed to 145GB/s since P4 would fail to reach the threshold of 165GB/s in certain scenarios.
- ▶ Fixed an issue with the targeted power test on T4 that would cause incorrect failures in some cases
- ▶ Fixed an issue with NVVS to report failures on a per-GPU basis
- ▶ Fixed an issue with NVVS to report failures on a per-GPU basis
- ▶ `dcgmFieldValue_t` is no longer supported in DCGM. The return value of the `dcgmGetLatestValuesForFields()` and `dcgmEntityGetLatestValues()` APIs is an updated struct `dcgmFieldValue_v1`, so developers may need to update their application to use the new struct when calling these APIs
- ▶ On K80s, failures due to throttling are disabled by default. See the Known Issues for more information

- ▶ Fixed issues with debug log file (`--debugLogFile`) and plugin statistics (`--statspath`) file generation with DCGM Diagnostics
- ▶ Fixed output formatting issues with `dcgmi diag --verbose`
- ▶ DCGM installer packages (deb and rpm) are now signed
- ▶ Fixed an issue with DCGM Diagnostics where in some cases, fields with the same timestamps are repeated in the statistics cache (available via log files)
- ▶ Fixed a limitation with the length of the log file name (specified using `debugLogFile`). The log file name including path can now support up to 128 characters

Known Issues

- ▶ When using profiling metrics with T4 in GPU VM passthrough, DCGM may report memory bandwidth utilization to be 12% higher.
- ▶ When using multiplexing of profiling metrics, the PCIe bandwidth numbers returned by DCGM may be incorrect. This issue will be fixed in a later release of the profiling metrics feature.
- ▶ On DGX-2/HGX-2 systems, ensure that `nv-hostengine` and the Fabric Manager service are started before using `dcgmproftester` for testing the new profiling metrics. See the Getting Started section in the DCGM User Guide for details on installation.
- ▶ 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.
- ▶ To report NVLINK bandwidth utilization DCGM programs counters in the HW to extract the desired information. It is currently possible for certain other tools a user might run, including `nvprof`, to change these settings after DCGM monitoring begins. In such a situation DCGM may subsequently return errors or invalid values for the NVLINK metrics. There is currently no way within DCGM to prevent other tools from modifying this shared configuration. Once the interfering tool is done a user of DCGM can repair the reporting by running `nvidia-smi nvlink -sc 0bz; nvidia-smi nvlink -sc 1bz`.

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.