



# DATACENTER GPU MANAGER 1.5.6

v1.5.6 | November 2018

## Release Notes



# TABLE OF CONTENTS

- Changelog..... iii
- New Features..... iii
- Improvements..... iii
- Bug Fixes..... iv
- Known Issues..... v

# CHANGELOG

## New Features

This version of DCGM (v1.5.6) 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.

### General

- ▶ DCGM supports a new modular architecture (see `dcgmi` modules and documentation for more information).
- ▶ Added APIs (`dcgmGetNvLinkLinkStatus`) and CLI options (`dcgmi nvlink -s`) to retrieve the NVLink link status for every GPU and NVSwitch in a system.
- ▶ Added a new script (`blacklist_recommendations.py`) to test for error conditions and recommend blacklisting of unhealthy GPUs in an NVSwitch system.
- ▶ Added `dcgmEntitiesGetLatestValues` as an API to get the latest value for a list of entities for a list of fieldIds. This API also allows the user to choose whether those values are live from the driver or retrieved from the cache..

### Platform Support

- ▶ Added support for NVSwitch based systems - NVIDIA DGX-2 and HGX-2 products via the NVIDIA fabric manager.
- ▶ Added support for the NVIDIA Tesla T4 product. See Known Issues for details on DCGM Diag (NVVS).
- ▶ Added support for the NVIDIA Tesla V100 32GB SKU.
- ▶ Added support for the NVIDIA Tesla P6 product.
- ▶ Added support for the NVIDIA Tesla M10 product.

## Improvements

### General

- ▶ Changed `dcgmGetDeviceAttributes` to use live value snapshots so that background watches aren't necessary between the host engine and the cache manager, thereby removing ~10 default field watches.

- ▶ Added a new throttle reasons field to DCGM. Some DCGM diag tests can now generate warning due to clock throttling.
- ▶ Memory bandwidth test is now supported on Tesla P4, Tesla K80 GPUs.
- ▶ Added a new `dcgmi` command line option to make output available in json format.
- ▶ DCGM diags now warns when single bit errors (SBEs) occur during a test rather than fail. Added options that allow the user to specify the SBE failure threshold.
- ▶ DCGM libraries are now installed under `/usr/lib/x86_64-linux-gnu` on Ubuntu systems.
- ▶ Added a sample config file for DCGM diag (under the `sdk_samples` directory).
- ▶ Added a sample to showcase blacklisting of DCGM modules (`dcgmModuleGetStatuses`, `dcgmModuleBlacklist`).
- ▶ Added support for additional Tesla V100 SKUs.
- ▶ Removed incompatible fields `DCGM_FI_DEV_COMPUTE_PIDS`, `DCGM_FI_DRIVER_VERSION`, `DCGM_FI_NVML_VERSION`, `DCGM_FI_PROCESS_NAME` from the collectd plugin.
- ▶ On ppc64 Ubuntu platform, DCGM libraries are installed in `/usr/lib/powerpc64le-linux-gnu`.
- ▶ DCGM documentation can now be found online at <http://docs.nvidia.com/datacenter/dcgm> and packages no longer include documentation.
- ▶ Changed DCGM diag to warn rather than fail when GPUs have thermal violations.
- ▶ Added checks to ensure the power health tests read the most recent power levels from the system.
- ▶ Added a number of improvements to `blacklist_recommendations.py`.

## Bug Fixes

- ▶ Fixed an issue where `dcgmi diag` would not read config files passed via the command line options.
- ▶ Fixed an issue where `dcgmi diag` CLI options were not consistent with the standalone `nvvs` tool.
- ▶ Fixed an issue where `dcgmi dmon` would not print requested metrics correctly on stdout.
- ▶ Fixed an issue where requesting only level 1 tests of DCGM diag would erroneously print that higher levels of tests were successful.
- ▶ Fixed an issue where DCGM diag (NVVS) would read incorrect device attributes on systems with heterogeneous GPUs.
- ▶ Fixed an issue in DCGM Diag (NVVS) to resolve errors such as "**Unable to retrieve pcie\_replay\_counter summary data: Bad parameter passed to function**" would be returned.
- ▶ Fixed an issue on ppc64 Ubuntu platforms where libraries were not loaded from the correct paths.
- ▶ Fixed an issue where DCGM diag (NVVS) would crash after listing the available GPUs with `-g`.

- ▶ Fixed an issue where DCGM diag would indicate success when run on an empty group (with no GPUs assigned).
- ▶ DCGM now correctly detects clock throttling on the GPU(s) where it happens rather than report it for each GPUs in a multi-GPU system.
- ▶ Fixed an issue where DCGM diag (NVVS) would timeout waiting for GPUs to return to 65C between test runs. This allows customers to let GPUs run as long as they are below max operating temperature/slowdown temperature.

## Known Issues

- ▶ DCGM Diag (NVVS) is does not support the NVIDIA Tesla T4 product. Support will be added in a future update of DCGM.

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.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

NVIDIA and the NVIDIA logo 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.

© 2013-2018 NVIDIA Corporation. All rights reserved.