



# DATACENTER GPU MANAGER 2.3

v2.3 | October 2021

## Release Notes



# TABLE OF CONTENTS

- Changelog..... iii
- Patch Releases..... iii
  - DCGM v2.3.5..... iii
  - DCGM v2.3.4..... iii
  - DCGM v2.3.2..... iv
- DCGM v2.3 GA..... iv
  - New Features..... iv
  - Improvements..... iv
  - Bug Fixes..... iv
  - Known Issues..... v

# CHANGELOG

This version of [DCGM](#) (v2.3) requires a minimum R418 driver that can be downloaded from [NVIDIA Drivers](#). On NVSwitch based systems such as DGX A100 or HGX A100, a minimum of Linux R450 ( $\geq 450.80.02$ ) driver is required. If using the new profiling metrics capabilities in DCGM, then a minimum of Linux R418 ( $\geq 418.87.01$ ) driver is required. It is recommended to install the latest datacenter driver from NVIDIA drivers downloads site for use with DCGM.

## Patch Releases

### DCGM v2.3.5

DCGM v2.3.5 released in March 2022.

#### Overview

- ▶ This release is a rebuild of DCGM 2.3.4 packages. No new features, improvements or bug fixes have been included in this patch release.

### DCGM v2.3.4

DCGM v2.3.4 released in February 2022.

#### Bug Fixes

- ▶ Fixed an issue where DCGM reported 0 for the memory usage (framebuffer) for MIG devices (on A100 and A30).
- ▶ Fixed an issue where DCGM would report 0s for profiling metrics when multiple processes would context switch on the GPU.
- ▶ Added missing NVLink error counters (CRC, flit, data, replay and recovery) when reported via `dcgmi nvlink --errors`
- ▶ Fixed an issue with DCGM diagnostics when running stress tests that target Tensor Cores on the GPUs.
- ▶ Security updates: See Security Bulletin for CVE-2022-21820: NVIDIA DCGM March 2022, on the [NVIDIA Product Security page](#).

## DCGM v2.3.2

DCGM v2.3.2 released in January 2022.

### Improvements

- ▶ Added support for the NVIDIA A2 product.

### Bug Fixes

- ▶ Fixed an issue with MIG support, where whole device metrics such as temperature, power etc. were reported as 0 for GPU instances.

## DCGM v2.3 GA

DCGM v2.3.1 released in October 2021.

## New Features

### General

- ▶ DCGM Diagnostics now accepts a configuration file to customize the thresholds specified for each GPU product. The main DCGM package (**datacenter-gpu-manager**) now has a dependency on a new config package (**datacenter-gpu-manager-config**) that can be updated independent of the main package.
- ▶ Added the ability to test for correctness of peer-to-peer (P2P) copies over the PCIe protocol (reads/writes, initiated by the device to host and data strides) in DCGM Diagnostics.
- ▶ Added support for NVIDIA A16 and RTX A5000/A4000 products.
- ▶ Added alerting on NVSwitch recovery and fatal errors using NVIDIA NSCQ.

## Improvements

- ▶ DCGM now includes a collectd **types.db** file for all DCGM fields. Refer to the User Guide for more information on using DCGM with collectd.
- ▶ Reduced the default CPU usage of DCGM-Exporter by reducing the watch interval of metrics.

## Bug Fixes

- ▶ Fixed incorrect error message reporting (number of seconds and samples) for thermal violations.
- ▶ Fixed an issue where DCGM was not reporting new NVIDIA brand types from NVML (when observed using **DCGM\_FI\_DEV\_BRAND**).
- ▶ Fixed an issue where **dsgm diag --debugLogFile** had an artificial character limit on file names.
- ▶ Fixed bug that could lead to infinite storage of metrics with sub-second polling intervals.

- ▶ Fixed Arm64 (aarch64) support for profiling metrics.

## Known Issues

- ▶ 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

© 2013-2022 NVIDIA Corporation. All rights reserved.