



NVIDIA DGX SuperPOD

Release Notes

Featuring NVIDIA DGX A100 and DGX H100 Systems

RN-11287-001 V12
2023-12-12
BCM 10.23.12

Contents

Chapter 1.	Introduction.....	1
Chapter 2.	Component Versions.....	2
Chapter 3.	Change Requests.....	3
	General3	
	CMDaemon.....	3
	Node Installer.....	4
	cm-diagnose.....	4
	cm-harbor.....	5
	cm-kubernetes-setup.....	5
	cm-scale.....	5
	cm-setup.....	5
	cm-wlm-setup.....	6
	jupyter 6	
	pythoncm.....	6

Chapter 1. Introduction

This document covers the NVIDIA Base Command™ Manager (BCM) 10.23.12 software release on NVIDIA DGX SuperPOD™ configurations. Except for Chapter 2, the information herein is the same as in the [*NVIDIA Base Command Manager Release Notes*](#).

Information about BCM and DGX SuperPOD is available at:

- > <https://docs.nvidia.com/base-command-manager/>
- > <https://docs.nvidia.com/dgx-superpod/index.html>

Chapter 2. Component Versions

DGX SuperPOD component versions for this release are in Table 1.

Table 1. Common component versions

Component	Version
BCM ISO	10.23.12
DGX OS	6.1.0
Ubuntu	Ubuntu 22.04.1 LTS
Enroot	3.4.1-1
CUDA toolkit	12.2
DCGM	3.1.8
Cumulus OS	5.5.1
Mellanox InfiniBand Switch (DGX H100)	MLNX OS version: 3.11.1014 HCA Firmware: CX7 - 28.36.2050
Mellanox InfiniBand Switch (DGX A100)	MLNX OS version: 3.11.1014 HCA Firmware: CX7 - 28.36.2050
Slurm	23.02.6
Mellanox OFED Driver (A100 and H100)	23.10-0.5.5.0 (Slogin and DGX nodes)
DGX kernel	5.15.0-1040-nvidia
GPU Driver	535.129.03
Lustre Client	lustre-client-modules-5.19.0-45-generic
UFM	UFM 3.0 SDN version: 1.3.1
HPL	hpc-benchmarks:23.10
NCCL	tensorrt:23.11-py3

Chapter 3. Change Requests

General

3.1.1 New Features

- > Added support for Kubernetes v1.28

3.1.2 Improvements

- > Added mlnx-ofed23.10 package
- > Added CUDA 12.3 packages
- > Update cuda-driver to 545.23.08
- > Update cm-openssl to 3.1.4
- > Update cuda-driver-legacy-470 to 470.223.02

CMDaemon

3.1.3 New Features

- > Allow the Kubernetes kubelet service to start when swap is enabled, which resolves an issue where Kubernetes setup may fail if the head node has swap and is selected for Kubernetes master node

3.1.4 Improvements

- > Added support for two VLANs on top of a bond interface with a VLAN as the provisioning interface
- > Added a periodic CMDaemon maintenance task to free the heap allocations that are not automatically freed, which reduces the memory usage
- > Display both the nodes and the accelerator counts in the license information in cmsh
- > Added an hourly cron job to clean up files left behind when the sample-ipmi script is being killed
- > Prevent already outdated monitoring data with timestamps in the past to be saved in the CMDaemon database

- > Added an “export as CSV” option in the REST monitoring data AP
- > Added new monitoring metrics ManagedServicesOK for the partition and the categories to get totals for all of the nodes in the cluster or the nodes in a given category
- > Allow the option to disable the nvlink and nvswitch metrics by specifying extra_values configuration settings for the nodes
- > Include the outdated packages for the passive head node, the node-installer, and the software images in the notification for available BCM updates

3.1.5 Fixed Issues

- > An issue with sorting on timestamps in a monitoring plot consisting of raw and consolidated data, which in some cases can result in CMDaemon returning no monitoring data for certain requests
- > An issue with reporting the correct InfiniBand count in the hardware overview of the compute nodes
- > In some cases, a timing issue that may prevent the pbsmom service from starting in an on-perm+edge workload manager setup
- > An issue with the Prometheus exporter for monitoring data when the data contains measurable names with spaces
- > An issue with performing parallel device power operations with cmsd
- > Allow the option to use GET in the Prometheus sampler for older exporters that do not allow POST operations
- > An issue with moving the software image revisions directories when updating the path of the parent software image

Node Installer

3.1.6 Improvements

- > On Ubuntu based distributions, make the updates of the ntp.conf drift file settings consistent between the Node Installer and CMDaemon, which until now could generate different configuration files

cm-diagnose

3.1.7 Improvements

- > Include syslog in cm-diagnose
- > Sanitize all mysqldumps in cm-diagnose

cm-harbor

3.1.8 Fixed Issues

- > In some cases, fixed a race condition where Harbor from the cm-harbor package and Shorewall are concurrently updating the iptables rules, which can prevent enabling the required iptables rules

cm-kubernetes-setup

3.1.9 Fixed Issues

- > An issue selecting the correct Kubernetes namespace in the retry mechanism of cm-kubernetes-setup when uninstalling failed operators

cm-scale

3.1.10 New Features

- > The Auto Scaler now takes into account the NVIDIA GPU requests made by Kubernetes pods and jobs when selecting the compute nodes to power on

3.1.11 Fixed Issues

- > Auto Scaler now takes Slurm mincpus parameter into account

cm-setup

3.1.12 Fixed Issues

- > An issue with ssh connections initiated by cm-setup scripts when a ssh-agent is running and has a SSH ECDSA certificate added to the agent
- > A regression in cm-container-registry-setup for Harbor on HA head node setup, which can result in "no such file or directory" error messages for files not present on the passive head node

cm-wlm-setup

3.1.13 Improvements

- > Create the enroot cache shared directory automatically if it does not already exist

3.1.14 Fixed Issues

- > An issue where enroot is not configured by default on a head node when the pyxis Slurm plugin is enabled
- > Allow the option to configure full BCM GPU autodetection for Slurm with cm-wlm-setup

jupyter

3.1.15 Improvements

- > In some cases, an issue where duplicated pods or services may be created due to a race condition in the Kubernetes API
- > Update the JupyterLab and JupyterHub dependencies to the most recent versions

pythoncm

3.1.16 Improvements

- > Added an option to use gzip compression in RPC calls

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.

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.

Trademarks

NVIDIA, the NVIDIA logo, NVIDIA Base Command, NVIDIA DGX, and NVIDIA DGX SuperPOD are trademarks and/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

© 2023 NVIDIA Corporation. All rights reserved.