



# NVIDIA Base Command Manager

## Release Notes

RN-17026-001 V3  
2023-10-23  
BCM 10.23.10

# Contents

Chapter 1.	Introduction .....	1
Chapter 2.	Change Requests .....	2
2.1	General .....	2
2.2	CMDaemon.....	2
2.3	cm-scale.....	3
2.4	Cloud.....	3
2.5	Kubernetes .....	4
2.6	Workload Management .....	4
2.7	Container Engines.....	5
2.8	Monitoring.....	5
2.9	Cluster on Demand .....	5
2.10	Base View .....	5
2.11	CMDaemon.....	6

# Chapter 1. Introduction

This release covers NVIDIA Base Command™ Manager (BCM) 10.23.10.

Previously release notes are located at:

<https://docs.nvidia.com/base-command-manager/#release-notes>

Information about BCM is available at:

<https://docs.nvidia.com/base-command-manager/>

>

# Chapter 2. Change Requests

## 2.1 General

### 2.1.1 New Features

- > Added mlnx-ofed23.07 package
- > Added cm-pmix4 package

### 2.1.2 Improvements

- > Added drainstatus to cm-diagnose
- > Updated cuda-driver package to 535.104.12
- > Updated cm-libprometheus package to 0.47.0
- > Updated cm-openssl package to 3.1.3

## 2.2 CMDaemon

### 2.2.1 New Features

- > Added advanced config flag DisableRemoteShell to disable all remote shell RPC
- > Added events for Cumulus service management operations

### 2.2.2 Improvements

- > Added cmsh clone device option to increment IP addresses by values other than 1
- > Allow lite node IP to be set during cmsh device add
- > Display an error when setting an invalid software image in cmsh
- > Update /etc/resolv.conf via netconfig on SLES15 instead of writing file
- > Created the ability to add model/serial number information to new switches (ZTP)
- > Kill active ramdisk create process when software image is removed

## 2.2.3 Fixed Issues

- > Fixed provisioning trigger when an image name starts with the name of another image
- > Allow `cm-cmd-ports --get` to work without an active `cmd`
- > Prevent "Reboot required: Interfaces have been modified" event from being shown for a node if the node has a VLAN interface on a Bridge interface that includes a bond interface
- > Fixed `cm-burn` unsuccessful completion in the absence of both a pre and post section
- > Image updates on provisioning nodes now wait for provisioning operations on other nodes to complete before proceeding.
- > Allow appending or skipping adding a Slurm drain reason when healthcheck fails with drain action enabled
- > Fixed crash of `pythoncm` parallel node termination function
- > Fixed an edge case that causes hostlist generation failures when there are 3 numeric fields in the hostname
- > Fixed service management for `cm-lite-daemon`

## 2.3 cm-scale

### 2.3.1 Fixed Issues

- > Allow to start terminated cloud nodes whose state is one of the node installer ones
- > Terminate useless AWS spot instance requests
- > Fixed the termination of cloud nodes when multiple clone operations are issued in parallel
- > Fixed the startup of nodes by `cm-scale` if Slurm job predicted start time is set by Slurm in the future
- > Fixed handling of job arrays with range from 1 to >1 figure number

## 2.4 Cloud

### 2.4.1 New Features

- > Added support for AWS FSx on Ubuntu for `cmjob`

### 2.4.2 Improvements

- > Improved error message when starting a cloud node with incorrect VPC/subnet configuration

## 2.4.3 Fixed Issues

- > Fixed issue with cm-cloud-storage-setup when using us-east-1 region
- > Prevent cloud instance termination when cloud director is down from being listed as UP+terminated
- > Fixed starting spot instances after a no-capacity in availability zone scenario occurs
- > Unfulfilled spot instance requests stay in PENDING state until fulfilled or terminated
- > Store availability zones for networks created by COD or manually, which enables AutoScaler to distribute loads between availability zones in COD deployments

## 2.5 Kubernetes

### 2.5.1 New Features

- > Added support for NGC token authentication in cm-kubernetes-setup

### 2.5.2 Improvements

- > Improved the wizard when it should fail earlier than it actually does (incorrect return code checks caused the installer to confusingly fail at later stages)
- > Kubernetes wizard errors will now show more context information where possible
- > Increased timeouts for kubeadm init and clusterctl init operations to effectively handle slow connections

### 2.5.3 Fixed Issues

- > Add user wizard will use BCM user name and not commonName

## 2.6 Workload Management

### 2.6.1 New Features

- > Added enroot and enroot+caps packages

### 2.6.2 Fixed Issues

- > Update AWS spot instances state in Slurm when they are terminated outside of BCM

## 2.7 Container Engines

### 2.7.1 Improvements

- > Improved internal IP detection logic for etcd (similarly to internal IP detection for Kubernetes Calico and Flannel)

## 2.8 Monitoring

### 2.8.1 New Features

- > Added Prometheus /rules and /alert and /alertmanagers end points
- > Added operstate metrics (operational state i.e., UP / DOWN ) via cm-lite-daemon for Cumulus switches

### 2.8.2 Improvements

- > Display K/M/G in cmsh for consolidated averages when no unit is set for a metric

### 2.8.3 Fixed Issues

- > Added support to run healthcheck with storcli software next to megacli software

## 2.9 Cluster on Demand

### 2.9.1 Improvements

- > Improved the display of the EULA when running from docker image
- > Allow CMDaemon to work with cluster-on-demand cluster spanning multiple regions (requires manual setup)

## 2.10 Base View

### 2.10.1 Improvements

- > Provide notifications in Base View if BCM package updates are available
- > Visualize licensed GPU used and available in Base View

## 2.11 CMDaemon

### 2.11.1 Improvements

- > Patched cm-curl package to address CVE-2023-38545



## 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.