



NVIDIA Base Command Manager

Release Notes

RN-17026-001 V1
2023-09-21
BCM 10.23.09

Contents

Chapter 1.	Introduction	1
Chapter 2.	Change Requests	2
2.1	General	2
2.2	CMDaemon.....	3
2.3	Node Installer	5
2.4	cm-scale.....	5
2.5	Linux and Hardware Integration.....	5
2.6	Cloud.....	6
2.7	Kubernetes	6
2.8	Improvements	6
2.9	Workload Management	7
2.10	Jupyter	7
2.11	Container Engines.....	8
2.12	Container Registries	8
2.13	Ceph.....	8
2.14	Monitoring.....	8

Chapter 1. Introduction

NVIDIA Base Command™ Manager (BCM) 10.23.09 is the first public release for version 10, a new major version of NVIDIA cluster management software.

Starting with this version, Bright Cluster Manager is being merged into Base Command Manager.

Base Command Manager 10 is the next release of Bright Cluster Manager 9.2, now under the new name.

Unless expressly stated otherwise, Base Command Manager 10 includes all the functionality supported by Bright Cluster Manager 9.2.

This version marks the first release of BCM also used with NVIDIA AI Enterprise. From the same codebase as BCM, Base Command Manager Essentials (BCME) will be packaged with features certified for the NVIDIA AI Enterprise use case.

BCM 10 is licensed per GPU. This is different from the legacy Bright Cluster Manager product, which was licensed per node. Customers with active support subscriptions using Bright Cluster Manager 9.2 and earlier can upgrade to BCM 10 by exchanging their current licenses for GPU-based BCM 10 licenses at no cost. Contact sw-bright-sales-ops@NVIDIA.onmicrosoft.com for more information about licensing.

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

- > Support for Oracle Cloud Infrastructure for Cluster On Demand
- > Support for NVIDIA Spectrum switches provisioning (Cumulus OS 5) and management via cm-lite-daemon
- > Support for NVIDIA BlueField-2 and BlueField-3 Data Processing Units (DPUs) provisioning (BFB) and management
- > Support for NVIDIA AI Enterprise software versions
- > New DGX SuperPOD post install setup tool (cm-pod-setup)
- > New DGX SuperPOD network configuration setup tool (bcm-netautogen)
- > Switch to GPU-based licensing
- > Add cm-cron service
- > Add cm-list-image-conf-files.py script to list all special files in <image>/cm/conf/
- > Add cuda12.2 packages
- > Add mlnx-ofed23.04 package
- > Add cuda-driver-legacy-470 package to support older datacenter/Tesla GPUs requiring NVIDIA CUDA driver version 470

2.1.2 Improvements

- > Update cm-openssl package to 3.1.2
- > Update mlnx-ofed58 package to 5.8-3.0.7.0
- > Update mpls-ofed54 package to 5.4-3.7.5.0
- > Update mlnx-ofed49 package to 4.9-7.1.0.0
- > Update mlnx-ofed59 DGX H100 package to 5.9.0.5.6.0.125

2.2 CMDaemon

2.2.1 New Features

- > Add `cmsh device switchports` command to get an overview of available switch ports
- > Send a warning event when a provisioning request has stalled longer than 2 hours. (Default value can be configured)

2.2.2 Improvements

- > Switch to UUIDs to uniquely identify entities
- > Allow `cm-mig-manage` to support GPUs that do not have `index = minorID`
- > Turn on MIG on DGX H100 after node reboot when `MIG.profiles` are set in GPU settings
- > Increase DHCP maximal search domains to 32 by default
- > Add `cmsh chassis set members` as compact device list
- > Preserve files in `/cm/images/<image>/cm/conf/{node,category}/` while updating images with `rsync`
- > Show an error message when `cmsh createramdisk` is run without arguments or an image set
- > Improved daily cron script to create monthly backup files for the `openldap-servers` to also include backups older than 1 year
- > Add a new `'--all'` option to `cmsh sysinfo` command to show extra information that has been collected by CMDaemon
- > Prevent CMDaemon crash when missing or truncated files are present in the monitoring backup directory
- > Increase `systemd-resolved.service` reload timeout
- > Redirect all `stdout/stderr` from a `cmburn` test script to a log file
- > Show inherited kernel properties in `cmsh device get`
- > Add multiline support for `cmsh rack display`
- > Add `free extra_values` to all entities to store additional information
- > Remove field for the CPU frequency scaling governor
- > Add `--certificate --key` options in `cmsh help`
- > Add `user/group` name validation in `cmsh`
- > Do not populate status for each node in the environment to avoid multiple slow RPCs

2.2.3 Fixed Issues

- > Fix killing jobs on a node when CMDaemon is restarted on that node
- > Fix RemoteMountChecker when a custom port is specified as the NFSCheckerPort AdvancedConfig parameter when querying cm-nfs-checker
- > Handle cm-lite-daemon restart properly
- > Fix help of cmsh cert removerequest command
- > Fix HPL test start in cmburn on SLES 15 base distribution
- > Automatically adjust overlay.category references when a category is removed
- > Do not clone switchports when cloning a device
- > Fix CMDaemon crash when malformed JSON data is sent
- > Update node environment cache when automatically changing FS exports
- > Honor backup role disabled=yes configuration
- > Detect xvd* disk in sysinfo
- > Prevent the addition of duplicate nameservers in /etc/resolv.conf
- > Delete duplicate entries in /etc/nginx/nginx.conf
- > Fix cmsh crash when cloning an entity without specifying a name in the genericresources submode
- > Hide all events in cmsh if --hide-events is used
- > Remove verbose logs in /tmp/aws* from cm-setup
- > Fix cmsh table formatting with long lines
- > Fix default gateway for edge nodes running Ubuntu
- > Fix duplicate nodes for monitoring pickup scheduler
- > Fix database storage of drained provisioning nodes
- > Ensure named gets reloaded when network changes made
- > Fix false negative open --failbeforedown when a status value is unchanged
- > Fix typo guage -> gauge

2.3 Node Installer

2.3.1 Fixed Issues

- > Fix booting of compute nodes with separate /usr filesystem
- > Allowed cloning of headnodes with btrfs filesystems
- > Fix disk management script to correctly assemble MD raids

2.4 cm-scale

2.4.1 New Features

- > Support for Oracle Cloud Infrastructure for Auto Scaler
- > Automatically detect memory and GPUs for cloud nodes

2.4.2 Improvements

- > Support multi-partition Slurm jobs in Auto Scaler

2.4.3 Fixed Issues

- > Fix incorrect number of CPUs for Slurm jobs in Auto Scaler
- > Handle lack of availability zone capacity for AWS spot instances in Auto Scaler
- > Auto Scaler ignores queue priorities for multi-queue Slurm jobs

2.5 Linux and Hardware Integration

2.5.1 New Features

- > Support for DGX OS 6.1
- > Add cm-dpu-setup tool to define NVIDIA BlueField-2 and BlueField-3 Data Processing Units (DPUs) in the cluster
- > Add cm-dpu-manage to perform management actions on NVIDIA BlueField-2 and BlueField-3 Data Processing Units (DPUs)

2.6 Cloud

2.6.1 New Features

- > Add cm-cod-oci to create Cluster on Demand in Oracle Cloud Infrastructure
- > Allow COD-AWS cluster to span multiple regions (Contact support for assistance)
- > Add support for AWS FSx on Ubuntu

2.6.2 Fixed Issues

- > Fix various issues with Azure locations caused by Azure API errors
- > Improved support for AWS spot instances

2.7 Kubernetes

2.7.1 New Features

- > Change Kubernetes deployment to use kubeadm
- > Change Kubernetes deployment to use packages from kubernetes.io instead of cm-kubernetesXXX packages
- > Support for Cluster API (CAPI) as a deployment method for new Kubernetes clusters

2.8 Improvements

- > Update Kyverno to 3.0.4 (due to incompatibility with Kubernetes 1.27.x)
- > Support for multiple NVIDIA GPU operator versions
- > Deploy the NVIDIA GPU Operator with toolkit.enabled=false by default

Fixed Issues

- * NVIDIA GPU Operator deployment always results in NVIDIA packages being installed
- * Update exclude lists for Kubernetes to avoid failures on "grabimage"
- * Do not include kubelet.service file in exclude list (this can interfere with assigning additional nodes to the Kubernetes roles and prevent the kubelet service from starting correctly)

2.9 Workload Management

2.9.1 New Features

- > Support data and cache sharing options for pyxis and enroot
- > Allow management of Slurm prolog/epilog timeouts

2.9.2 Improvements

- > Rely on MIG autodetection to configure gres.conf
- > Update Slurm package to 23.02 (older versions are not supported anymore)
- > Use pmix4 with Slurm 23.02
- > pyxis may now be compiled and installed from a local tarball with sources
- > All RPCs for job management API in CMDaemon also return an exit code of the operation

2.9.3 Fixed Issues

- > Fix parsing of Slurm job CPUs
- > Fix fetching job information when UGE accounting rotation is configured
- > Fix UGE AdditionalSubmitHosts advanced configuration flag
- > Advanced accounting (job types and account hierarchy monitoring)

2.10 Jupyter

2.10.1 New Features

- > Manage Spark and PostgreSQL instances from JupyterLab
- > Manage Pods and data migration from/to Persistent Volume Claims
- > Read Pod logs and events from Jupyter interface
- > Support for multi-factor authentication

2.10.2 Improvements

- > Support for private NGC credentials in Kubernetes kernel templates

2.11 Container Engines

2.11.1 Improvements

- > Update cm-docker package to 23.0.6
- > Update cm-containerd package to 1.7.1
- > Update cm-apptainer package to 1.1.9

2.12 Container Registries

2.12.1 Improvements

- > Update cm-harbor package to 2.8.2
- > Update cm-docker-registry package to 2.8.1

2.12.2 Fixed Issues

- > Generate containerd certificates when a registry mirror is not configured

2.13 Ceph

2.13.1 Improvements

- > Updated Ceph to Ceph Quincy

2.14 Monitoring

2.14.1 New features

- > Add new NVSwitch metrics
- > Support for Graphana 10

2.14.2 Improvements

- > Disable job metrics collection when JobSampler is not setup to run in OOB mode
- > Sample node JobsRunning metric even when there are no jobs running
- > Reduce memory usage spike when using PromQL over short timespans
- > Multiply metric value by 100 when displaying % in pythoncm
- > Exclude rdma* by default in /proc/net/dev sampler
- > Exclude virtual ibp*v* interface from monitoring

2.14.3 Fixed Issues

- > Fix the Slurm job_gpu_utilization and job_gpu_wasted metric calculations when running GPU process within sbatch scripts
- > Fix calculation of job_gpu_wasted metric when the node has multiple GPUs
- > Fix samplenow CPUUsage metric
- > Ensure job_gpu_* have correct values in the first few seconds of a job being started
- > Ensure first data sample of a Prometheus sampler is stored to the database
- > Propagate cumulative values passed by a JSON sampler during initialize
- > Fix metrics sampling when temperatures are not provided by the Redfish API
- > Clean up job monitoring when jobs are removed from cache

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.