



NVIDIA DGX SuperPOD

Release Notes

RN-11287-001 V15
2024-05-24
BCM 10.24.05

Contents

Chapter 1.	Introduction.....	1
Chapter 2.	Component Versions.....	2
Chapter 3.	Change Requests.....	3
3.1	General.....	3
3.2	CMDaemon.....	3
3.3	Base View.....	5
3.4	Cluster Tools.....	5
3.5	COD.....	6
3.6	Machine Learning.....	6
3.7	cm-clone-install.....	6
3.8	cm-cluster-extension.....	6
3.9	cm-create-image.....	7
3.10	cm-kubernetes-setup.....	7
3.11	cm-scale.....	7
3.12	cm-wlm-setup.....	8
3.13	cmsh.....	8
3.14	jupyter.....	8
3.15	pythoncm.....	9

Chapter 1. Introduction

This document covers the NVIDIA Base Command™ Manager (BCM) 10.24.05 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/>

Chapter 2. Component Versions

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

Table 1. Common component versions

Component	Version
BCM ISO	10.24.05
DGX OS	6.2.1
Ubuntu	22.04.2 LTS
Enroot	3.4.1
CUDA toolkit	12.2
DCGM	3.3.5
Cumulus OS	5.5.1
Mellanox InfiniBand Switch (DGX H100)	MLNX OS version: MLNX-OS 3.11.2016 HCA Firmware: CX7 - 28.39.2048
Mellanox InfiniBand Switch (DGX A100)	MLNX OS version: MLNX-OS 3.11.2016 HCA Firmware: CX6 - 20.39.2048
Slurm	23.02.7
Mellanox OFED Driver (A100 and H100)	MLNX_OFED_LINUX-23.10-2.1.3.1 LTS
DGX kernel	5.15.0-1053-nvidia
GPU Driver	535.161.08
Lustre Client	lustre-client-modules-5.19.0-45-generic 2.14.0-ddn125
UFM	NVIDIA UFM Enterprise Appliance Software v1.7.0
HPL	nvcr.io/nvidia/hpc-benchmarks :24.03
NCCL	nvcr.io/nvidia/tensorrt :24.02-py3
DGX FW	1.1.3

Chapter 3. Change Requests

3.1 General

3.1.1 New Feature

- > Added mlnx-ofed24.01 package
- > Added CUDA 12.4 toolkit packages
- > Added PBS Professional 2024 packages
- > cmdaemon-apidocs has been replaced with cm-api-docs; Documentation can be accessed via the landing page
- > Added cm-nsight-systems-cli package containing the latest CLI version of nsight-systems package, replacing the current cm-nsight-systems package
- > *bcm-post-install: Allow slogin node base name to configurable*
- > *bcm-post-install: added command for SuperPOD validation*

3.1.2 Improvements

- > Updated mlnx-ofed23.10 to 23.10-2.1.3.1
- > For new Ubuntu 22.04 head node installations, use fixed port numbers for the NFS lockd, statd, and mountd daemon

3.1.3 Fixed Issues

- > An issue that prevents cm-diagnose from completing when single quotes are used in cmd.conf for dbuser/dbpass
- > An issue with the runtime and PID path settings in the nvidia-persistenced service unit file from the cuda-driver—packages
- > *Fix SuperPOD enroot configuration*

3.2 CMDaemon

3.2.1 New Features

- > Mark the devices monitored by MQTT as UP when recent monitoring data exists

- > Allow the option to run Slurm accounting database daemon in high-availability mode

3.2.2 Improvements

- > Added per node Slurm state metric
- > Allow the option to select automatically a random free port for the IMEX service
- > Allow the option to define an exclude list in the network interfaces healthcheck to skip specified interfaces
- > Added `/var/lib/rancher/.*` to the default exclude list for the ProcMounts monitoring producer, which otherwise can create unnecessary monitoring metrics
- > Improved performance of the cm-mqtt service
- > Allow the option to automatically set the bond and the bond members MAC addresses in the CMDaemon node interfaces entities when the node boots
- > Include the `perm-mac-address` under a bond interface when verifying the license, which resolves an issue with verifying the license when a bond network interface is created after the license is requested
- > Added `lxc-` interfaces to the default exclude list for the ProcNetDev monitoring producer
- > Allow the option to disable a MQTT with a flag in the configuration file

3.2.3 Fixed Issues

- > An issue with validating of the LDAP group during commit of a user
- > An issue with clearing the memory when a large number of entities that have failing health checks are added and then removed
- > An issue where DNS allow-query configuration entries are not added on edge directors for Kubernetes networks, preventing queries from these networks
- > An issue where monitoring consolidators may not be created for all entity-measurable pairs
- > An issue where a temporary `resolv.conf` bind mount created in the software images are being added to the CMDaemon monitoring database
- > Added a retry mechanism around `gethostbyname` when writing the IMEX configuration files, which can otherwise throw an exception
- > An issue with chargeback calculations using per node requested CPU/GPU information
- > An issue in the drain action manager code that in some cases can lead to high CMDaemon memory usage
- > An issue with determining the number of requested CPUs for multi-node Slurm jobs when storing the jobs information in CMDaemon

- > An issue which can lead to high CMDaemon memory usage if the post-provisioning monitoring-resume operation has failed
- > An issue with hard-coded references to /sbin/arping which in some cases can prevent CMDaemon from using arping in the event of a failover
- > An issue in the RPC status code which in some cases can result in an infinite recursion on the passive head node in the event of a failover
- > An issue with the node profile missing the UPDATE_CONFIG_FILES_AFTER_IMAGE_UPDATE_TOKEN which prevents the /cm/conf files from being copied from the software image to the provisioned nodes when using non-head-node provisioners
- > An issue with the reporting of the GPU chargeback information
- > An issue where a category can be removed while another category's provisioning role still has a reference to it
- > An issue where Azure cloud compute nodes are cloned with an incorrect power status when the original node's power status is ON
- > An issue where on failure AWS node power on actions produce an error message "Unable to parse output" instead of the AWS error message
- > An issue with the cmsh dropunused command that can result in removing too many measurables
- > An issue with the cmsh device syncinfo command when specifying an fs part path

3.3 Base View

3.3.1 Fixed Issues

- > An issue with displaying the SNMP system information data for switches

3.4 Cluster Tools

3.4.1 Fixed Issues

- > An issue where cm-mysql-sanitize.py, which is required by cm-diagnose, is not part of the cluster-tools package

3.5 COD

3.5.1 New Features

- > Added support for creating HA COD clusters in Azure
- > Allow the option to skip shared storage setup with the cm-cloud-ha-setup tool
- > Added support for OCI defined tags. This changes the original --head-node-tags command line option to --head-node-freeform-tags and adds new command line option --head-node-defined-tags

3.5.2 Improvements

- > Allow the option to select the Azure availability zone on the command line of the cluster create command

3.6 Machine Learning

3.6.1 New Features

- > Introduced ML NCCL and CuDNN packages for CUDA 12.4

3.6.2 Fixed Issues

- > An issue where WLM kernels may be unexpectedly restarted if one of the kernels fails to start

3.7 cm-clone-install

3.7.1 Fixed Issues

- > An issue with handling bond interfaces and bond members configuration on Ubuntu base distribution

3.8 cm-cluster-extension

3.8.1 Fixed Issues

- > An issue where 'germany' is incorrectly listed as an Azure region

3.9 cm-create-image

3.9.1 Fixed Issues

- > An issue where missing modular metadata for the 'default' package group on the RHEL8 and RHEL9 ISOs can prevent the creation of software images

3.10 cm-kubernetes-setup

3.10.1 Improvements

- > Ensure the `/var/lib/etcd` directory has the correct permissions (0700) for etcd member to be able to join the etcd cluster

3.10.2 Fixed Issues

- > A regression in `cm-kubernetes-setup` that allows the user to select nodes in a way that results in the overlap of compute nodes between different Kubernetes clusters
- > An issue with `cm-kubernetes-setup -pull` unable to complete if while the images are being pulled the pod is evicted due to disk pressure

3.11 cm-scale

3.11.1 New Features

- > Allow the option to reboot the nodes in FULL install mode when the `cm-scale` engine is changed

3.11.2 Fixed Issues

- > An issue where the shutdown state from files may be used incorrectly

3.12 cm-wlm-setup

3.12.1 Fixed Issues

- > Setting up pyxis will no longer configure it to clean the data directory from epilog since enroot can perform this automatically

3.13 cmsh

3.13.1 Improvements

- > Allow the options to specify the IP increment with the cmsh addinterface command
- > Include the job run time data in the cmsh WLM jobs info command
- > Allow the option to specify the network CIDR on the cmsh “add network” command line

3.13.2 Fixed Issues

- > An issue where the cmsh monitoring trigger info command does not show grouped expressions
- > An issue with importing older formats of the .cmshhistory file, which can result in duplicating all entries in the cmsh command history

3.14 jupyter

3.14.1 New Features

- > Allow the option to use sqsh files to run Jupyter kernels based on enroot
- > Restrict the access to Jupyter based on group memberships

3.14.2 Improvements

- > Allow the option to install and configure VNC when setting up Jupyter

3.15 pythoncm

3.15.1 Fixed Issues

- > An issue in the pythoncm cluster.py implementation where an incorrect logger variable name is being used
- > pyxis-sources

3.15.2 Improvements

- > Updated pyxis-sources to 0.19.0

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.