



# NVIDIA DGX SuperPOD

## Release Notes

RN-11287-001 V16  
2024-07-29  
BCM 10.24.07

# 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	Node Installer.....	5
3.4	Cluster Tools.....	6
3.5	COD.....	6
3.6	Head Node Installer.....	7
3.7	Machine Learning.....	7
3.8	cm-docker-setup.....	7
3.9	cm-kubernetes-setup.....	7
3.10	cm-lite-daemon.....	8
3.11	cm-scale.....	8
3.12	cmsh.....	8
3.13	jupyter.....	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	10.24.07
DGX OS	6.2.1
Ubuntu	22.04. <a href="#">24</a> LTS
Enroot	3.5.0
CUDA toolkit	12.2
DCGM	3.3.5
Cumulus OS	5.5.1
Mellanox InfiniBand Switch (DGX A100/H100)	MLNX OS version: MLNX-OS 3.11.2016 HCA Firmware: CX7 - 28.39.2048
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	ddn145
UFM	NVIDIA UFM Enterprise Appliance Software v1.7.0
HPL	nvcr.io/nvidia/hpc-benchmarks:23.10
NCCL	nvcr.io/nvidia/tensorrt:24.06-py3
DGX FW	1.1.3

# Chapter 3. Change Requests

## 3.1 General

### 3.1.1 New Feature

- > Added support for air-gapped Kubernetes setup
- > Added support for Slurm 24.05
- > Added support for Kubernetes v1.29

### 3.1.2 Improvements

- > Updated Ubuntu 22.04 base distribution to 22.04.4
- > Updated cm-nvhpc to 24.5
- > Updated cm-openssl to 3.1.6
- > Updated cuda-driver-535 to 535.183.06
- > Updated DCGM to 3.3.6
- > Updated enroot to 3.5.0
- > Updated Lmod to 8.7.39
- > Updated PBS Professional 2022 to 2022.1.6
- > Updated PBS Professional 2024 to 2024.1.1
- > Updated Slurm 23.11 to 23.11.8

## 3.2 CMDaemon

### 3.2.1 New Features

- > Allow the option to configure the real hostname of the head node instead of the “master” hostname alias as the StorageHost in slurmdbd.conf
- > Use the hardware information which corresponds to the cloud instance type when configuring cloud compute nodes in Slurm. This allows the nodes to be configured before the nodes have been powered on for the first time
- > Allow the option to update the IB switches firmware via cmsh or pythoncm
- > Added support for accelerated networking (SR-IOV) in Azure

- > Allow the option to use NVIDIA sharp plugin as the Slurm resource selection plugin (SelectType)
- > Added a new OOM kill count metric based on data from /proc/vmstat
- > Added new GPU/NVlink monitoring metrics for CRC errors and remapped rows
- > Added cm-remove-orphaned-pending-job-info.py helper script that can be used to remove jobs information in the CMDaemon database for WLM jobs that have been cached by cmd as pending while they are removed from the workload manager
- > Allow the option with Ubuntu base distributions to configure a layer3 network setup with a /31 connection between the individual node switch port and the router
- > Spread out the schedulers health check to reduce the flood of parallel WLM calls
- > Display the standalone entities as a path in the monitoring tree in Base View
- > Include the version information in the info-message reported by the cuda-dcgm health check when the health check is sampled with the -debug option
- > Allow the option to select the ICMP protocol for the firewall role's openports options
- > The DNS allow-transfer setting is now set to 'none' by default for all zones maintained by CMDaemon
- > Allow the option to define extra Prometheus labels for devices
- > Allow the option to expose enum values as labels with the Prometheus exporter
- > Allow the option to dump the monitoring data for health checks and enum metrics grouped by value
- > Allow the option to update the "FrozenFile" setting with the cm-manipulate-advanced-config.py helper script
- > Improved execution speed of the rogueprocess health check
- > Improved speed of the monitoring triggers evaluation when a regex is used
- > Validate the project managers when adding, updating, or removing users
- > Added a commit validation warning when a GPU sampler update frequency is too high. Too high frequency can lead to incorrect values for certain GPU metrics
- > Improved IMEX epilog stop script with a timeout of 15s and ability to kill the process
- > Allow the option to perform a periodic check if the head node IP has been changed on external DHCP renewal
- > Allow the option to change the behavior of the monitoring drain action to not set a drain reason when draining the node(s)

### 3.2.2 Fixed Issues

- > An issue in pythoncm prometheus range\_query not converting the interval to nanoseconds
- > An issue with using cmha status for failover groups

- > An issue with freezing the `/etc/systemd/resolved.conf` on Ubuntu base distributions
- > An issue with the `sample_ibmetrics.py` script not returning floating point values
- > An issue with adding the `additionalHostnames` of `GenericDevice` and `LiteNode` devices to the DNS configuration
- > An issue with applying node tags to bare-metal instances in OCI
- > An issue with providing Kubernetes job information to the Auto Scaler
- > An issue with the interfaces health check when the interface's speed is defined with a unit
- > An issue that prevents `auditd` from being added or removed when SELinux settings are changed for a partition or a category
- > An issue that can cause `CMDaemon` to crash on clusters with head node HA setup when an RPC is called within a small window during startup
- > An issue with setting the user and group ownership of static configuration files managed by generic roles
- > An issue where executing MIG, BIOS, or DPU commands may not clear the "busy" flag when the commands time out
- > Allow the option to select the AWS node-installer volume type with an advanced configuration option `NodeInstallerEbsVolumeType`
- > An issue where the values of enum metrics may not be translated to enums on the compute nodes until `CMDaemon` is restarted
- > Standardize all monitoring scripts to use `CMD_SCRIPTTIMEOUT` environment variable. The `CMD_SCRIPT_TIMEOUT` environment variable is no longer passed to the monitoring action scripts, now only `CMD_SCRIPTTIMEOUT` is used
- > An issue with the monitoring trigger actuator when many samples are picked up at once or when using complex expressions
- > An issue where a Cumulus ZTP script may not use the hostname provided by DHCP
- > An issue with generating `Slurm topology.conf` in the case when the switches are connected to both nodes and other switches
- > An issue where `CMDaemon` may execute `systemctl daemon-reload` also when the Slurm service drop-in file has not changed
- > An issue with creating AWS compute nodes with multiple EFA interfaces
- > An issue with parsing of the requested CPUs setting for UGE jobs information

## 3.3 Node Installer

### 3.3.1 Fixed Issues

- > An issue with setting up bonded network interfaces on diskless nodes

- > An issue where the auditd service may not be disabled by the node-installer when SELinux is not enabled

## 3.4 Cluster Tools

### 3.4.1 Improvements

- > Allow the option to create clusters spanning over multiple existing AWS subnets
- > Allow the option to setup Azure HA COD clusters by using existing public IP resources
- > Allow the option to not create a shared public IP for HA COD clusters in OCI
- > Allow the option to deploy docker on diskless nodes

## 3.5 COD

### 3.5.1 Improvements

- > The /etc/cm-install-release information file is now created also for cluster on demand head nodes
- > Updated oci-hpc-network-device-names to 1.0.13 with support for L40S
- > Allow the option to create clusters spanning over multiple existing AWS subnets
- > Added support for creating AWS clusters in a subnet outside of the main VPC CIDR block

### 3.5.2 Fixed Issues

- > An issue with creating COD clusters when FIDO/U2F keys (e.g. ED25519-SK) are added to a running SSH agent:
- > An issue where the delete ^^dry-run output does not show the actual resources that are to be deleted
- > An issue with restricting SSH access to OCI clusters to only a predefined CIDR due to default OCI security list rules
- > An issue with creating COD clusters when encrypted SSH keys are used



## 3.6 Head Node Installer

### 3.6.1 Improvements

- > Updated the STIG disk setup configurations

## 3.7 Machine Learning

### 3.7.1 Improvements

- > Updated CuDNN to 9.1.1

## 3.8 cm-docker-setup

### 3.8.1 Improvements

- > Allow the option to deploy docker on diskless nodes

## 3.9 cm-kubernetes-setup

### 3.9.1 New Features

- > Added support for GPU operator v24.3.0
- > Added support for Kubernetes v1.29
- > Updated local path provisioner to 0.0.28

### 3.9.2 Improvements

- > The Grafana web interface is now added to the cluster landing page for new Kubernetes setup installations when the Prometheus Operator stack is selected in cm-kubernetes-setup
- > Use Kyverno 1.11.4 for new Kubernetes deployments and by default add tolerations to allow running Kyverno on the control-plane nodes, and not only on the worker nodes, which allows the worker nodes to be shutdown without affecting Kyverno
- >

## 3.9.3 Deprecated Features

- > Kubernetes v1.24 and v1.25 are no longer available options for setting up Kubernetes

## 3.10 cm-lite-daemon

### 3.10.1 Fixed Issues

- > An issue with Cumulus switch overview showing state=unknown when CumulusOS 5.9 is installed

## 3.11 cm-scale

### 3.11.1 Fixed Issues

- > In some cases, an issue with saving the state of drained nodes when the head node is restarted, which can prevent the Auto Scaler from considering the nodes as available for scaling up

## 3.12 cmsh

### 3.12.1 Improvements

- > Allow the option with cmsh to set a trigger expression as a parsed string
- > Allow the option to show raw numbers and columns with all zero data in the chargeback report table

### 3.12.2 Fixed Issues

- > An issue with the cmsh addinterface command assigning incremental IPs without sorting the nodes by their hostnames
- > An issue that can cause cmsh to crash when running many cmsh -c '...' commands one after the other

## 3.13 jupyter

### 3.13.1 Improvements

- > Migrate the WLM kernel templates to Jupyter Kernel Starter

## 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.