



NVIDIA DGX SuperPOD

Release Notes

Featuring NVIDIA DGX A100 and DGX H100 Systems

RN-11287-001 V13
2024-02-13
BCM 10.24.01

Contents

Chapter 1.	Introduction.....	1
Chapter 2.	Component Versions.....	2
Chapter 3.	Change Requests.....	3
	General	
	CMDaemon.....	3
	cm-kubernetes-setup.....	4
	cm-lite-daemon.....	5
	cm-wlm-setup.....	5
	cmsh	5
	pythoncm.....	5

Chapter 1. Introduction

This document covers the NVIDIA Base Command™ Manager (BCM) 10.24.01 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	10.24.01
DGX OS	6.1.0
Ubuntu	Ubuntu 22.04.2 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: 3.11.2016 HCA Firmware: CX7 - 28.39.2048
Mellanox InfiniBand Switch (DGX A100)	MLNX OS version: 3.11.2016HCA Firmware: CX7 - 28.39.2048
Slurm	23.02.7
Mellanox OFED Driver (A100 and H100)	23.10-1.1.9.0 LTS (Slogin and DGX nodes)
DGX kernel	5.15.0-1042-nvidia
GPU Driver	535.129.03
Lustre Client	lustre-client-modules-5.19.0-45-generic 2.14.0-ddn125
UFM	UFM Enterprise SW: 6.15.1-4
HPL	hpc-benchmarks:23.10
NCCL	tensorrt:23.12-py3
DGX FW	1.1.3

Chapter 3. Change Requests

General

3.1.1 New Feature

- > The head node installer will now create a new `/etc/cm-install-release` file to keep a record of the installation time and media that has been used
- > Added support for upgrading BCM3/Bright9.2 clusters to BCM 10

3.1.2 Improvements

- > Added `cuda-driver-535` package
- > The `mlnx-ofed` packages' installation scripts will now pin down the kernel packages for Ubuntu when deploying MOFED
- > Updated `mlnx-ofed58` to 5.8-4.1.5.0
- > Updated `mlnx-ofed23.10` to 23.10-1.1.9.0
- > Updated `cuda12.3` to 12.3 update 2
- > Updated `cm-nvhpc` to 23.11

CMDaemon

3.1.3 Improvements

- > Update the Kubernetes users' configuration files with `Run:ai` configuration settings
- > Redirect the output from `cm-burn` to `tty1`
- > Added new GPU totals metrics for temperature and `nvlink` bandwidth
- > Allow the option to select BCM GPU autodetection configuration mechanism also in the Slurm WLM cluster settings, and not only in the Slurm WLM client role
- > Ensure kubelets are able to join a Kubernetes cluster also after the initial certificates have expired (which typically happens after 4 hours)

3.1.4 Fixed Issues

- > An issue with sorting the data passed to the PromQL engine, which can result in an error "expanding series: closed SeriesSet" when running instant queries
- > An issue where the exclude list snippets are not being cloned when cloning a software image

- > Rare deadlock in CMDaemon which can occur while committing a head node
- > An issue with configuring the GATEWAYDEV on RHEL 9 when a VLAN network interface is configured on top of the BOOTIF interface
- > An issue where `/etc/systemd/resolved.conf` is not being added to the imageupdate exclude list for the compute nodes
- > An issue where `install-license` may not copy some certificates to all `/cm/shared*` on a multi-arch or multi-os cluster
- > An issue with the Prometheus exporter when entities have recently been removed
- > An issue with parsing multiple pending Kubernetes CSR per node, which can result in none of the CSR's being approved
- > On SLES base distribution, an issue with updating the cluster landing page with links to the dashboards of other integrations such as Kubernetes or Ceph
- > An issue where CMDaemon may not restart the Slurm services automatically when the number of CPUs configuration settings for some nodes change
- > An issue where CMDaemon may hang waiting for events while stopping
- > An issue where the `cmsh` call to create a certificate may return before the certificate is written
- > An issue where entering the `cmsh biosettings` mode may result in an "Error parsing JSON" error message
- > In some cases, an issue with configuring Slurm when GPU automatic configuration by BCM has been selected
- > In some cases, an issue with setting up Etcd due to insufficient permissions to access the Etcd certificate files
- > An issue where a WLM job process ID may be added to an incorrect cgroup, which in some cases may result in the process being killed when another WLM job running on the same node completes
- > An issue with collecting GPU job metrics for containerized Pyxis jobs

cm-kubernetes-setup

3.1.5 New Features

- > Use Calico 3.27, Run:ai 2.15.2, and GPU operator v23.9.1 for new Kubernetes deployments using `cm-kubernetes-setup`

3.1.6 Improvements

- > Allow the option to choose Network Operator version 23.10.0
- > Allow the option to configure a custom Kubernetes Ingress certificate

3.1.7 Fixed Issues

- > An issue with setting up Run:ai which can result in the Run:ai cluster installer not being able to complete successfully
- > An issue with the interactive uninstall question in cm-kubernetes-setup when the Kubernetes API is not responsive
- > In some cases, an issue where cm-kubernetes-setup may not wait for the required nodes, such as control-plane or worker nodes, to come back up after a reboot

cm-lite-daemon

3.1.8 Improvements

- > Added new metrics for the total traffic on network interfaces

cm-wlm-setup

3.1.9 Fixed Issues

- > In some cases, an issue with installing Pyxis on multi-arch or multi-distro software images
- > Pyxis enroot is now configured to use its internal value for the cache directory, which previously was being set to a directory under /run

cmsh

3.1.10 New Features

- > Added a new cmsh “multiplexers” command in monitoring / setup which can show which nodes will run for other entities a specified dataproducer

pythoncm

3.1.11 Improvements

- > Added a new pythoncm example script total-job-power-usage.py for calculating WLM jobs power usage

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.