



NVIDIA DGX SuperPOD

Release Notes

RN-11287-001 V14
2024-03-27
BCM 10.24.03

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	Head Node Installer.....	5
3.5	cm-kubernetes-setup.....	6
3.6	cm-scale.....	6
3.7	cm-wlm-setup.....	6
3.8	cmsh.....	6
3.9	pyxis-sources.....	7
3.10	Slurm23.11.....	7

Chapter 1. Introduction

This document covers the NVIDIA Base Command™ Manager (BCM) 10.24.03 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	10.24.03
DGX OS	6.2.0
Ubuntu	22.04.3 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 - 20.39.1002
Slurm	23.02.7
Mellanox OFED Driver (A100 and H100)	23.10-1.1.9.0 LTS
DGX kernel	5.15.0-1046-nvidia
GPU Driver	535.161.07
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 :23.10
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 support for RHEL8u9 and ROCKY8u9
- > Added support for RHEL9u3 and ROCKY9u3
- > The DGX OS software images included on the head node installer ISOs are now based on DGX OS 6.2.0 (Release 1)
- > Enable NVSM metrics for DGX systems
- > Added a new BCM package for NVIDIA Nsight systems (cm-nsight-systems)

3.1.2 Improvements

- > Include the gsp firmware with the cuda-driver- packages
- > Updated DCGM to 3.3.5
- > Updated cuda-driver to 550.54.15
- > Updated cuda-driver-535 to 535.161.07
- > Updated cuda-driver-legacy-470 to 470.239.06
- > Updated munge to v0.5.16
- > Updated cm-nvhpc to 24.1
- > Updated cm-openssl to 3.1.5
- > Removed the dependency of the cm-nvidia-container-toolkit package on the cuda-driver package, which otherwise can cause some package managers to remove the toolkit package when the CUDA driver is replaced with a different version

3.1.3 Fixed Issues

- > Increase the stack and nofile limits in cm-config-limits for the root user on Ubuntu 22.04 to prevent possible issues with head nodes hanging under heavy load

3.2 CMDaemon

3.2.1 New Features

- > An issue where CMDaemon may store the Slurm array's main job information in the CMDaemon monitoring DB, which creates unnecessary entries in the DB since these jobs expand to individual tasks as they get are scheduled

- > Allow the option to open a remote-request-assistance session from within cmsh
- > CMDaemon user profiles can now include new tokens, such as SET_USER_PROFILE_TOKEN, allowing a cmsh or Base View client connected using a certificate with this profile to set or update the profile setting for users
- > CMDaemon will now generate a combined kubeconfig file in .kube/config in the home directory of a user containing all clusters the user has access to, allowing the user to connect to the Kubernetes cluster without first loading the environment module
- > Allow the option to use negative matching such as “!resource!=category-name” in the monitoring comparison expressions

3.2.2 Improvements

- > Improved Prometheus exporter cache management to ensure the memory usage does not grow over time
- > CMDaemon will now set the InfiniBand (IB) interface GUID as extra values for the IB interfaces, which then can be shown in cmsh / Base View
- > Exclude the link MAC entry from the Cumulus switch overview information
- > Allow the option to set the MAC address of non-node devices to the MAC address identified on the respective switch port for the device
- > Restrict users’ abilities with a profile that allows them to add other users to also set arbitrary group ids for the users they create. The behavior can be tuned with advanced configuration options
- > Increase the Slurm job queue QOS table size in the CMDaemon DB
- > Redirect all base-view and userportal HTTPS calls from the passive to the active head node
- > CMDaemon will now send an event when network interfaces are added or removed

3.2.3 Fixed Issues

- > An issue where setting the extra_values nvlink property to false is not picked up by CMDaemon, causing CMDaemon to continue sampling nvlink metrics
- > An issue that can result in duplicate entries in the /exporter Prometheus endpoint
- > An issue where a network switch ZTP.sh can point to an incorrect IP address for the head node
- > An issue where a disabled backup role is still being used for monitoring backup, preventing it from being removed from the target list
- > An issue where a small buffer size for the user’s groups can prevent CMDaemon from correctly listing all user’s groups
- > Ensure the head nodes can provision each other regardless of the ProvisioningRole selections
- > In some cases, an issue with folding the compute nodes hostnames when generating slurm.conf configuration file

- > An issue with Prometheus sampler with exporter that only supports http GET
- > An issue where committing a monitoring action without setting a script can cause CMDaemon on the head node to crash
- > An issue with Slurm job management operations in cmsh and Base View unable to handle Slurm job arrays ids
- > An issue where metrics from the AggregateNode producer do not have correct data / do not have “no data” values when there are no nodes in a rack
- > In some cases, an issue that can leave a passive head node CMDaemon process using 100% CPU
- > In some cases, an issue that may prevent CMDaemon from loading old jobs information from the CMDaemon DB
- > An issue with sending a test email when using cmsh
- > An issue with the passive head node forwarding labeled entity information to the active, preventing it from being used in PromQL queries
- > An issue where terminated cloud nodes that are consequently powered on may still be listed as ‘terminated’ in cmsh
- > An issue that prevents two different configuration overlays with the same priority and different generic roles to be committed in CMDaemon

3.3 Node Installer

3.3.1 Fixed Issues

- > An issue where setting the frozenFilesPerNode directive may not cause the node-installer to freeze /etc/sysconfig/network on RHEL

3.4 Head Node Installer

3.4.1 New Features

- > Updated the default partition sizes for the standard RAID1 and RAID5 head node disk layouts to match the sizes of the non-raid standard layout

3.4.2 Fixed Issues

- > An issue with head node installations with Lmod where the DefaultModules.lua module file is not created by default, resulting in messages about empty LMOD_SYSTEM_DEFAULT_MODULES environment variable

3.5 cm-kubernetes-setup

3.5.1 Fixed Issues

- > Improved error reporting when kubeadm init step fails

3.6 cm-scale

3.6.1 Improvements

- > Improved handling of a lack of vCPUs in AWS in the same way as a lack of capacity

3.7 cm-wlm-setup

3.7.1 New Features

- > Allow the option to set the enroot temporary directory using cm-wlm-setup

3.7.2 Fixed Issues

- > Ensure cm-wlm-setup can install AGE 2023.1.1 (8.8.1)

3.8 cmsh

3.8.1 Improvements

- > Added new cmsh WLM jobs mode command pidsgpus to list the pids and the gpus used by a WLM job

3.8.2 Fixed Issues

- > An issue in cmsh user mode with case-sensitive compare of profile names

3.9 pyxis-sources

3.9.1 New Features

- > Updated pyxis sources package to 0.17.0

3.10 Slurm23.11

3.10.1 Improvements

- > Updated slurm23.11 to 23.11.4

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.