



# NVIDIA DGX SuperPOD

## Release Notes

Featuring NVIDIA DGX A100 Systems

# Document History

RN-11287-001

<b>Version</b>	<b>Date</b>	<b>Authors</b>	<b>Description of Change</b>
01	2023-02-07	Robert Sohigian	NVIDIA Base Command™ Manager (BCM) release 3.23.01
02	2023-02-11	Rangam Addepalli and Robert Sohigian	Corrected component versions in Table 1
03	2023-02-22	Robert Sohigian	Changes to Table 2 and formatting updates

# Contents

1.	Introduction .....	1
2.	Component Versions .....	2
3.	DGX SuperPOD Documentation .....	3
4.	Change Requests .....	4
4.1	General .....	4
4.1.1	Improvements .....	4
4.1.2	Fixed Issues .....	4
4.2	CMDaemon .....	4
4.2.1	New Features .....	4
4.2.2	Improvements .....	4
4.2.3	Fixed Issues .....	5
4.3	Head Node Installer .....	5
4.3.1	Fixed Issues .....	5
4.4	Machine Learning .....	5
4.4.1	Improvements .....	5
4.5	cm-wlm-setup .....	6
4.5.1	New Features .....	6

---

# 1. Introduction

These release notes appertain to the NVIDIA Base Command Manager (BCM) 3.23.01 software release on the NVIDIA DGX SuperPOD™.

The document is in three sections:

- > Component Versions
- > DGX SuperPOD Documentation
- > Change Requests

---

## 2. Component Versions

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

**Table 1. Component versions**

Component	Version
Bright	9.2-8
BCM ISO	bcm3_installer-3.23.01-1.iso
Slurm	21.08.8-2
Enroot	3.4.0
Ubuntu	20.04.05
DGX-OS	DGX_SWBUILD_VERSION="5.4.0"
mlnx-ofed-all	MLNX_OFED_LINUX-5.4-3.4.0.0
GPU driver	520.141.03
Kernel	5.4.0-124-generic
that DCGM	2.4.5
CUDA Version	11.4
Cumulus OS	5.2.1
Lustre Client	2.14.0_ddn54

---

## 3. DGX SuperPOD Documentation

Links to DGX SuperPOD documentation are in Table 2.

NVIDIA Bright Cluster Manager documentation is available on [docs.nvidia.com](https://docs.nvidia.com).

**Table 2. DGX SuperPOD documentation**

Document	Link
NVIDIA DGX SuperPOD Administration Guide	<a href="#">GPU Genius</a>
NVIDIA DGX SuperPOD Data Center Design Guide	<a href="#">GPU Genius</a>
NVIDIA DGX SuperPOD Deployment Guide	<a href="#">GPU Genius</a>
NVIDIA DGX SuperPOD Reference Architecture	<a href="#">GPU Genius</a>
NVIDIA DGX SuperPOD Release Notes	<a href="#">GPU Genius</a>
NVIDIA DGX SuperPOD Solution Design	<a href="#">GPU Genius</a>
NVIDIA DGX SuperPOD User Guide	<a href="#">GPU Genius</a>
Cabling Data Centers	<a href="#">GPU Genius</a>

---

## 4. Change Requests

Change requests addressed in Bright Cluster Manager 9.2-8 are listed in this section.

### 4.1 General

#### 4.1.1 Improvements

- > Integration with Run:ai

#### 4.1.2 Fixed Issues

- > `mlnx-ofed57`: An issue where the `mst` service does not start when using `mlnx-ofed57` due to the `mst sysvinit` service file being packaged as a `systemd` unit file

### 4.2 CMDaemon

#### 4.2.1 New Features

- > Add a `cm-package-release-info` tool which can determine the Bright 9.X-Y version of installed packages

#### 4.2.2 Improvements

- > Introduce a new `CPUUsage` metric for compute and head nodes to show the percentage of the CPU usage
- > Exclude the ram and loop devices from `/sys/block/` from being sampled by the `SysBlockStat` monitoring data producer
- > Allow for special OID for PDU load to be specified via the revision property

## 4.2.3 Fixed Issues

- > Disable the `cgroup` job metrics collection for users (user IDs) that cannot be found on the nodes
- > An issue where the `monitoringdrop` command may drop the data only for the head node
- > An issue with creating the LSF configuration when some node is converted from a compute to a submit-only host
- > An issue where cloud node that have never been booted may have a status "unknown error"
- > Automatically start `slurmdbd` when Slurm configuration is frozen in `cmd.conf`
- > An issue where the lite daemon may not reconnect when the websocket is closed cleanly on the server side

## 4.3 Head Node Installer

### 4.3.1 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

## 4.4 Machine Learning

### 4.4.1 Improvements

Introduced packages:

- > `cm-onnx-pytorch-*-cuda11.7-*`
- > `cm-gpytorch-*-cuda11.7-*`
- > `cm-fastai2-*-cuda11.7-*`
- > `cm-pytorch-extra-*-cuda11.7-*`
- > `cm-xgboost-*-cuda11.7-*`
- > `cm-cub-cuda11.7`
- > `cm-tensorflow2-*-cuda11.7-*`
- > `cm-opencv4-*-cuda11.7-*`
- > `cm-ml-pythondeps-*-cuda11.7-*`
- > `cm-pytorch-cuda11.7`



#### Updated packages:

- > `cm-gcc9-*` to v9.5.0
- > `cm-pytorch-*` to v1.13.0
- > `cm-tensorflow2-*` to v2.10.0
- > `cm-gpytorch-*` to v1.9.0
- > `cm-fastai2-*` to v2.7.0
- > `cm-xgboost-*` to v1.6.2

#### Deprecated packages:

- > `cm-openmpi4-cuda11.2-ofed47-gcc9`
- > `cm-openmpi4-cuda11.2-ofed51-gcc9` packages

#### General:

- > Deprecated packages for CUDA 11.2 and introduced new variants for CUDA 11.7

## 4.5 `cm-wlm-setup`

### 4.5.1 New Features

- > Allow to use the "master" keyword in the `cm-wlm-setup` configuration as a placeholder for the real head node host name

## 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 NON-INFRINGEMENT, 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.