Hewlett Packard Enterprise

Technical white paper

# NVIDIA GPU Docker container installation guide for CentOS 7.6 support in CUDA 10.1 Update 2

# Contents



## Introduction

HPE has collaborated with NVIDIA®. to harnesses the superior parallel processing of NVIDIA GPUs with a comprehensive set of computing and infrastructure innovations. Using GPU-accelerated Docker containers allows users to take advantage of the ease of deployment and agility of Docker containers coupled with the parallel processing power of the programmable NVIDIA GPUs on select HPE servers.

### About this document

This document is intended for users familiar with Linux®. environments. This guide covers only the installation steps for enabling Docker containers to run on NVIDIA GPUs using CentOS 7.6® supporting the CUDA 10.1 Update 2 platform.

## System requirements

To use the NVIDIA Container Toolkit, you need a GPU that supports NVIDIA Compute Unified Device Architecture (CUDA)®.

**Important**
NVIDIA GPU drivers (part of the NVIDA CUDA Toolkit 10.1 Update 2 installation) must be installed. The containers will not run without the GPU drivers.
Superuser privileges are required for installation. Use the superuser (root) or use the sudo prefix if the sudo package is enabled.

NVIDIA CUDA Toolkit 10.1 Update 2 is available at: http://developer.nvidia.com/cuda-downloads

Requirements for the NVIDIA CUDA Toolkit on a CentOS 7.6 system are listed in Table 1.

**Table 1.** NVIDIA CUDA Toolkit 10.1 Update 2 requirements

| CentOS 7.6 | Requirements |
|---|---|
| **Kernel** | 3.10 |
| **GCC** | 4.8.5 |
| **GLIBC** | 2.17 |
| **ICC** | 19.0 |
| **PGI** | 18.x,19.z |
| **XLC** | NO |
| **CLANG** | 8.0.0 |

## Install Docker 19.03 or later

To deploy NVIDA GPU Docker containers, you must install Docker 19.03 or later. To install Docker, issue the following commands:

```
[root@gpu2 ~]# yum install -y yum-utils
[root@gpu2 ~]# yum-config-manager --add-repo https://download.docker.com/linux/centos/docker-
ce.repo
[root@gpu2 ~]# yum install -y docker-ce
[root@gpu2 ~]# systemctl start docker
[root@gpu2 ~]# groupadd docker
[root@gpu2 ~]# gpasswd -a ${USER} docker
[root@gpu2 ~]# systemctl restart docker
[root@gpu2 ~]# newgrp docker
```

### (Optional) Verify Docker

Verify that Docker is installed by issuing the following command:

```
[root@gpu2 ~]#  docker run hello-world
```

## Install the NVIDIA-Docker Container Toolkit

**Important**
If nvidia-docker 1.0 is installed, you must remove it and all existing GPU containers.

To uninstall nvidia-docker 1.0:

```
[root@gpu2 ~]# docker volume ls -q -f driver=nvidia-docker | xargs -r -I{} -n1 docker ps -q -a -f
volume={} | xargs -r docker rm -f
[root@gpu2 ~]# sudo yum remove nvidia-docker
```

To install the NVIDIA-GPU Docker Container Toolkit, you first need to add package repositories:

```
[root@gpu2 ~]# distribution=$(. /etc/os-release;echo $ID$VERSION_ID)
[root@gpu2 ~]# curl -s -L https://nvidia.github.io/nvidia-docker/$distribution/nvidia-docker.repo |
sudo tee /etc/yum.repos.d/nvidia-docker.repo
[root@gpu2 ~]# yum install -y nvidia-container-toolkit
[root@gpu2 ~]# systemctl restart docker
```

Then verify the nvidia-smi install with the latest official CUDA image:

```
[root@gpu2 ~]# docker run --gpus all nvidia/cuda:9.0-base nvidia-smi
```

Sample output:

```
Mon Sep 16 16:55:46 2019
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 418.87.00    Driver Version: 418.87.00    CUDA Version: 10.1     |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Quadro P4000        Off  | 00000000:37:00.0 Off |                  N/A |
| 49%   39C    P0    28W / 105W |      0MiB /  8119MiB |      0%      Default |
+-------------------------------+----------------------+----------------------+
|   1  Quadro P4000        Off  | 00000000:AF:00.0 Off |                  N/A |
| 50%   40C    P0    27W / 105W |      0MiB /  8119MiB |      0%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                       GPU Memory |
|  GPU       PID   Type   Process name                             Usage      |
|=============================================================================|
|  No running processes found                                                 |
+-----------------------------------------------------------------------------+
```

**Additional links**
https://github.com/NVIDIA/nvidia-docker

## Learn more at
https://www.hpe.com/us/en/solutions/hpc-high-performance-computing/nvidia-collaboration.html

Make the right purchase decision. Click here to chat with our presales specialists.

**Share now**

**Get updates**

© Copyright 2019 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Docker is a trademark or registered trademark of Docker, Inc. in the United States and/or other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. NVIDIA is a trademark and/or registered trademark of NVIDIA Corporation in the U.S. and other countries. All other third-party marks are property of their respective owners.

a00090048ENW, September 2019