



# NVIDIA DOCA

## Developer Quick Start Guide

# Table of Contents

Chapter 1. Introduction.....	1
Chapter 2. Install DPU.....	2
Chapter 3. Register to NVIDIA DOCA Developer Zone.....	3
Chapter 4. Install SDKM.....	4
Chapter 5. Install DOCA Software Package.....	5
Chapter 6. Access DPU.....	6
Chapter 7. Run Reference DOCA Application.....	7
Chapter 8. More Information.....	8

---

# Chapter 1. Introduction

DOCA brings together a wide range of powerful APIs, libraries, and frameworks for programming and accelerating modern data center infrastructures. Like NVIDIA® CUDA® for GPUs, DOCA is a consistent and essential resource across all existing and future generations of DPU products.

This document is intended for those wishing to develop applications using the DOCA framework.



Note: Not sure which installation type to use? To expand on different DOCA user types and the relevant installation for each, see [NVIDIA BlueField and DOCA User Types](#).

---

## Chapter 2. Install DPU

Install the DPU into your host according to the installation instructions in the [DPU's hardware user guide](#). The steps include installing the DPU into the PCIe slot and properly securing it in the chassis. Make sure your host OS is listed under the "Supported Operating Systems" section in the [NVIDIA DOCA Installation Guide for Linux](#).

---

## Chapter 3. Register to NVIDIA DOCA Developer Zone

As part of the installation process, you must provide your NVIDIA account information. Get your NVIDIA Developer Zone access [here](#).

---

## Chapter 4. Install SDKM

Download [NVIDIA SDK Manager \(SDKM\)](#) and install it on the machine hosting the DPU.

- ▶ For Ubuntu/Debian host, run:

```
host# sudo apt install ./sdkmanager_[version]-[build]_amd64.deb
```

- ▶ For RedHat/CentOS host, run:

```
host# sudo yuminstall ./sdkmanager_[version]-[build#].x86_64.rpm
```

---

# Chapter 5. Install DOCA Software Package

SDKM installs all required development software on both the host and DPU. The following command installs or upgrades the required software on the DPU and host, flashes the new OS image and firmware, configures a network interface between the host and the DPU (RShim), and more.

Run this command on the host as a non-root user:

- ▶ For BlueField-2:

```
host# sdkmanager --cli install --logintype devzone --product DOCA --version latest --targetos Linux --host --target BLUEFIELD2_DPU_TARGETS --flash all
```

- ▶ For BlueField-3:

```
host# sdkmanager --cli install --logintype devzone --product DOCA --version latest --targetos Linux --host --target BLUEFIELD3_DPU_BMC_100GBE --flash all
```

HOST COMPONENTS	STATUS
▶ DOCA Software Package	[██████████] Ⓞ Installing - 25%
▶ Development Container on Host	[██████████] Ⓞ Install Pending

  

TARGET COMPONENTS	STATUS
▶ BlueField OS	[██████████] Ⓞ Flash Pending
▶ Firmware and Flashing Tools	[██████████] Ⓞ Pending
▶ Burn Target Firmware	[██████████] Ⓞ Install Pending
▶ DOCA Software Package	[██████████] Ⓞ Install Pending
▶ Development Container on Target	[██████████] Ⓞ Install Pending

During installation, you must perform the following steps:

- ▶ Provide root password to the host machine
- ▶ Set a new password to access the BlueField DPU (default username: `ubuntu`)



Note: DOCA software can also be installed manually using `apt-get` via the standard Linux Debian Package Manager. See the [NVIDIA DOCA Installation Guide for Linux](#) for more.

---

## Chapter 6. Access DPU

After a successful installation, you may log into the DPU's Ubuntu-based OS by running the following command from the host:

```
host# ssh ubuntu@192.168.100.2
```

Use the DPU password you defined during the installation process.

At this stage DOCA is installed on the DPU and the host server.



---

# Chapter 7. Run Reference DOCA Application

DOCA package assets (e.g., references, tools) are located on the DPU and on the host under `/opt/mellanox/doca/`.

The DOCA package includes a set of reference applications to facilitate developer onboarding. Please refer to the [NVIDIA DOCA Applications Overview](#) and [NVIDIA DOCA Programming Guides Overview](#) for more information.

To run the [DOCA File Scan](#) reference application which performs high-speed hardware RegEx acceleration using the RegEx engine:

1. Allocate huge pages and verify that the RegEx engine is active:

```
# sudo echo 2048 > /sys/kernel/mm/hugepages/hugepages-2048kB/nr_hugepages
# systemctl start mlx-regex
```

2. Compile the example RegEx rules file:

```
# rxcpc -f /opt/mellanox/doca/applications/file_scan/bin/regex_rules.txt -p 0.01 -o /tmp/regex_rules
```

3. Run the application using the following command:

```
# /opt/mellanox/doca/applications/file_scan/bin/doca_file_scan -p 03:00.0 -r /tmp/regex_rules.rof2.binary -d /opt/mellanox/doca/applications/file_scan/bin/data_to_scan.txt -l 60
Successful result: "Job 0 complete. Detected 7 match(es)"
```



Note: All parameters are available in the [NVIDIA DOCA File Scan Application Guide](#).

---

## Chapter 8. More Information

To learn more about NVIDIA BlueField DPUs, visit [NVIDIA BlueField DPU Documentation](#).

## 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 nor any of its direct or indirect subsidiaries and affiliates (collectively: "NVIDIA") make no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assume 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.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

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, and Mellanox are trademarks and/or registered trademarks of Mellanox Technologies Ltd. and/or NVIDIA Corporation in the U.S. and in other countries. The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a world-wide basis. Other company and product names may be trademarks of the respective companies with which they are associated.

## Copyright

© 2023 NVIDIA Corporation & affiliates. All rights reserved.