



NVIDIA DOCA Applications Overview

Guide

Table of Contents

Chapter 1. Introduction.....	1
1.1. Installation.....	1
1.2. Compilation.....	1
1.3. Developer Configurations.....	2
Chapter 2. Application Use of DOCA Libs.....	3
Chapter 3. Applications.....	5
3.1. Allreduce.....	5
3.2. App Shield Agent.....	5
3.3. Application Recognition.....	5
3.4. DMA Copy.....	5
3.5. DNS Filter.....	6
3.6. DPA All-to-all.....	6
3.7. DPA L2 Reflector.....	6
3.8. East-West Overlay Encryption.....	6
3.9. File Compression.....	6
3.10. File Integrity.....	6
3.11. File Scan.....	7
3.12. Firewall.....	7
3.13. GPU Packet Processing.....	7
3.14. Intrusion Prevention System.....	7
3.15. IPsec Security Gateway.....	7
3.16. NAT.....	7
3.17. OVS L4 Firewall.....	8
3.18. Programmable Congestion Control.....	8
3.19. Secure Channel.....	8
3.20. Simple Forward VNF.....	8
3.21. Switch.....	8
3.22. URL Filter.....	8
3.23. YARA Inspection.....	9

Chapter 1. Introduction

DOCA applications are an educational resource provided as a guide on how to program on the NVIDIA® BlueField® DPU using DOCA API.

For instructions regarding the development environment and installation, refer to the [NVIDIA DOCA Developer Guide](#) and the [NVIDIA DOCA Installation Guide for Linux](#) respectively.

1.1. Installation

DOCA applications are installed under `/opt/mellanox/doca/applications` with each application having its own dedicated folder. Each application installs two directories:

- ▶ `src` – contains the source code and compilation files
- ▶ `bin` – contains pre-built binaries alongside runtime configuration files

The application's binary is located under `/opt/mellanox/doca/applications/<application_name>/bin/doca_<application_name>`.

1.2. Compilation

As applications are shipped alongside their sources, developers may want to modify some of the code during their development process and then recompile the applications. The files required for the compilation are the following:

- ▶ `/opt/mellanox/doca/applications/meson.build` – main compilation file for a project that contains all the applications
- ▶ `/opt/mellanox/doca/applications/meson_options.txt` – configuration file for the compilation process
- ▶ `/opt/mellanox/doca/applications/<application_name>/src/meson.build` – application-specific compilation definitions

To recompile the application:

1. Move to the application's directory:

```
cd /opt/mellanox/doca/applications
```

2. On an environment that includes the [DPA All-to-all](#) application, ensure that a suitable environment variable is declared:

```
export PATH=/usr/mpi/gcc/openmpi-4.1.5rc2/bin:${PATH}
```

Refer to the [application's guide](#) for more information.

3. Prepare the compilation definitions:

```
meson /tmp/build
```

4. Compile the applications:

```
ninja -C /tmp/build
```

The generated applications will be located under the `/tmp/build/` directory.



Note: Compilation against DOCA's SDK relies on environment variables which are automatically defined per user session upon login. For more information, please refer to section "Meson Complains About Missing Dependencies" in the [NVIDIA DOCA Troubleshooting Guide](#).

1.3. Developer Configurations

When recompiling the applications, meson compiles them by default in "debug" mode. Therefore, the binaries would not be optimized for performance as they would include the debug symbol. For comparison, the application binaries shipped as part of DOCA's installation are compiled in "release" mode. To compile the applications in something other than debug, please consult [Meson's configuration guide](#).

The applications also offer developers the ability to use the DOCA Developer Log (`DOCA_DLOG`) on top of the existing DOCA Runtime Log (`DOCA_LOG`). Enabling the developer log during compilation activates various developer log messages that were left out of the release compilation. Activating the developer log may be done through `enable_developer_log` in the `meson_options.txt` file, or directly from the command line:

1. Prepare the compilation definitions to use the developer log:

```
meson /tmp/build -Denable_developer_log=true
```

2. Compile the application(s):

```
ninja -C /tmp/build
```

Chapter 2. Application Use of DOCA Libs

The following table maps DOCA reference applications to the libraries they make use of.

Application Category	BareMetal/Virtualized Cloud	Secure Cloud Gateway	Cloud Monitoring Storage	Streaming HPC														
Library	Flow	DPA	DMA	FlexIO	PCC	RegEx	App Shield	DPI	IPsec	SHA	Compression	Telemetry	GPU	Reverse	Network	GPU	Channel	
Network																		
DNS																		
Filter																		
GPU																		
Packet Processing																		
NAT																		
Simple																		
Forward																		
VNF																		
Switch																		
Security																		
App Shield Agent																		
Application Recognition																		
East-west Overlay Encryption																		
Firewall																		
Intrusion Prevention System																		
IPsec Security Gateway																		

	Open vSwitch				
	L4 Firewall				
	Secure Channel				#
	URL Filter			#	
	YARA Inspection			#	#
Data Path Acceleration	DPA all	#			
	DPA L2 Reflector		#		
	PCC			#	
Storage	DMA Copy		#		#
	File Compression			#	#
	File Integrity			#	#
	File Scan			#	
HPC	Allreduce				#

Chapter 3. Applications

3.1. Allreduce

[This application](#) is a collective operation that allows data from many processing units to be collected and merged into a global result before being delivered to all processing units using an operator. The application is implemented using the UCX communication framework, which leverages the DPU's low-latency and high-bandwidth utilization of its network engine.

3.2. App Shield Agent

[This application](#) describes how to build secure process monitoring and is based on the DOCA APSH library, which leverages DPU capabilities such as regular expression (RXP) acceleration engine, hardware-based DMA, and more.

3.3. Application Recognition

[This application](#) identifies applications that are in use on a monitored networking node. The application is based on the deep packet inspection (DPI) library, which leverages DPU capabilities such as regular expression (RXP) acceleration engine, hardware-based connection tracking, and more.

3.4. DMA Copy

[This application](#) describes how to transfer files between the DPU and the host. The application is based on the direct memory access (DMA) library, which leverages hardware acceleration for data copy for both local and remote memory.

3.5. DNS Filter

[This application](#) offloads DNS requests from the host to the DPU's Arm cores which allows reducing CPU overhead as they allow further DNS processing (e.g., allow/deny list) to be done. The application is based on the DOCA Flow and RegEx libraries which leverage DPU capabilities such as regular expression (RXP) acceleration engine, building generic execution pipes in HW, and more.

3.6. DPA All-to-all

[This application](#) is a collective operation that allows data to be copied between multiple processes. This application is implemented using DOCA DPA, which leverages the data path accelerator (DPA) inside of the BlueField-3 which offloads the copying of the data to the DPA and leaves the CPU free for other computations.

3.7. DPA L2 Reflector

[This application](#) uses the data path accelerator (DPA) engine to intercept network traffic and swap the source and destination MAC addresses of each packet. It is based on the FlexIO API which leverages DPU capabilities such as high-speed DPA.

3.8. East-West Overlay Encryption

[This application](#), IPsec, sets up encrypted connections between different devices and works by encrypting IP packets and authenticating the packets' originator. It is based on a strongSwan solution which is an open-source IPsec-based VPN solution.

3.9. File Compression

[This application](#) shows how to compress and decompress data using hardware acceleration and to send and receive it. The application is based on the DOCA Compress and DOCA Comm-Channel libraries.

3.10. File Integrity

[This application](#) shows how to send and receive files in a secure way using the hardware Crypto engine. It is based on the DOCA SHA and DOCA Comm Channel libraries.

3.11. File Scan

[This application](#) describes how to scan a file using the hardware RegEx engine to find whether there are matches according to the compiled regular expressions. It is based on the DOCA RegEx library which leverages the DPU's regular expression (RXP) acceleration engine.

3.12. Firewall

[This application](#) applies network security based on DOCA Flow gRPC and is used for remote programming of the DPU hardware. It leverages DPU capabilities such as building generic execution pipes in hardware, monitoring incoming and outgoing network traffic, and more.

3.13. GPU Packet Processing

[This application](#) shows how to combine DOCA GPUNetIO, DOCA Ethernet, and DOCA Flow to manage ICMP, UDP, TCP and HTTP connections with a GPU-centric approach using CUDA kernels without involving the CPU in the main data path.

3.14. Intrusion Prevention System

[This application](#) monitors a network for malicious activities or policy violations and is based on a deep packet inspection (DPI) library, which leverages DPU capabilities such as the regular expression (RXP) acceleration engine, hardware-based connection tracking, and more.

3.15. IPsec Security Gateway

[This application](#) demonstrates how to insert rules related to IPsec encryption and decryption based on the DOCA Flow and IPsec libraries, which leverage the DPU's hardware capability for secure network communication.

3.16. NAT

[This application](#), network address translation, switches packets with local IP addresses to global ones and vice versa. It is based on the DOCA Flow library which leverages DPU hardware capabilities such as building generic execution pipes in the hardware, executing specific actions on the traffic, and more.

3.17. OVS L4 Firewall

[This application](#) performs basic access control list (ACLs) operations. It allows the identification of different flows based on L3/L4 headers and executes different actions using Open vSwitch (OVS) commands.

3.18. Programmable Congestion Control

[This application](#), programmable congestion control, is based on the DOCA PCC library and allows users to design and implement their own congestion control algorithm, giving them good flexibility to work out an optimal solution to handle congestion in their clusters.

3.19. Secure Channel

[This application](#) is used to establish a secure, network-independent communication channel between the host and the DPU based on the DOCA Comm Channel library.

3.20. Simple Forward VNF

[This application](#) is a forwarding application that takes VXLAN traffic from a single RX port and transmits it on a single TX port. It is based on the DOCA Flow library which leverages DPU capabilities such as building generic execution pipes in the hardware, and more.

3.21. Switch

[This application](#) is used to establish internal switching between representor ports on the DPU. It is based on the the DOCA Flow library which leverages DPU capabilities such as building generic execution pipes in the hardware, and more.

3.22. URL Filter

[This application](#) limits access by comparing web traffic against a database to prevent users from different threats (e.g., malware, harmful sites, phishing). It is based on a deep packet inspection (DPI) library, which leverages DPU capabilities such as the regular expression (RXP) acceleration engine, hardware-based connection tracking, and more.

3.23. YARA Inspection

[This application](#) describes how to build YARA rule inspection for processes and is based on the DOCA APSP library, which leverages DPU capabilities such as the regular expression (RXP) acceleration engine, hardware-based DMA, and more.

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.