NVIDIA DOCA Applications Overview

Guide

MLNX-15-060539 _v1.4 | September 2022

Table of Contents

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 respectively.

Chapter 2. Applications

2.1. 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.

2.2. 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.

2.3. 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.

2.4. 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.

2.5. 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.

2.6. 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.

2.7. 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.

2.8. IPS

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.

2.9. L4 OVS 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.

2.10. 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.

2.11. 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 DOCA Flow which leverages DPU capabilities such as building generic execution pipes in the hardware, and more.

2.12. 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.

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 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, 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

NVIDIA Corporation | 2788 San Tomas Expressway, Santa Clara, CA 95051
http://www.nvidia.com