

NVIDIA DOCA Samples

Overview

MLNX-15-060528_v1.3 | May 2022

Table of Contents

Chapter 1. Introduction	1
Chapter 2. App Shield Samples	. 2
Chapter 3. Comm Channel Samples	3
Chapter 4. DPI Samples	. 4
Chapter 5. DMA Samples	. 5
Chapter 6. Flow Samples	. 6
Chapter 7. RegEx Samples	.7
Chapter 8. Telemetry Samples	8

Chapter 1. Introduction

The DOCA samples are simple code snippets meant to demonstrate the functionality of different DOCA lib programming on top of the DPU using the DOCA API.

Each sample demonstrates the use of only one DOCA library and focuses on a single component. The samples provide a reference for development and illustrate how to use the library API, but do not contain a main function or an executable to run. For a reference to a Meson build file, see the application's Meson file.

DOCA samples can be used either on the host or Arm, except for the App Shield samples which can only be used on the Arm side.

For instructions on the development environment and installation, refer to the <u>NVIDIA DOCA</u> <u>Developer Guide</u> and the <u>NVIDIA DOCA Installation Guide</u>.

Chapter 2. App Shield Samples

App Shield samples demonstrate how to monitor the host in several aspects.

The samples show how to properly initialize the DOCA App Shield lib and use its API to get information on the monitored system/process.

Chapter 3. Comm Channel Samples

Comm Channel samples demonstrate how to create a classic client-server framework. The client is created on the host side and a server on the DPU.

Chapter 4. DPI Samples

DPI samples demonstrate an end-to-end flow of a packet being inspected by the DOCA DPI library. The samples also show the use of gRPC over DPI library.

Chapter 5. DMA Samples

DMA samples demonstrate how to use the DMA library to copy memory from one buffer to another on the DPU. They also show how to copy memory from x86 host to the DPU.

Chapter 6. Flow Samples

Flow samples demonstrate how to create different types of pipes, add entries, and handle aged entries. The samples show the use of library structs (e.g., doca_flow_match, doca flow actions) for creating rules.

Chapter 7. RegEx Samples

RegEx samples demonstrate how to use the DOCA RegEx API to configure, send and receive data buffers to and from the BlueField hardware regular expressions#(RegEx)#engine, and get the matches according to the loaded compiled regular expression.

Chapter 8. Telemetry Samples

Telemetry samples demonstrate the use of the DOCA Telemetry and Telemetry NetFlow APIs. The samples show how to create a schema, start NetFlow, create source, and send records to a file or to IPC sockets.

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 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

© 2022 NVIDIA Corporation & affiliates. All rights reserved.

