



NVIDIA DOCA

Release Notes

Table of Contents

- Chapter 1. Introduction..... 1
 - 1.1. Installation Notes..... 1
 - 1.2. Technical Support..... 1
- Chapter 2. Changes and New Features..... 2
- Chapter 3. New Documentation..... 4
- Chapter 4. Known Issues..... 6

Chapter 1. Introduction

The NVIDIA® DOCA SDK is a software suite supporting the NVIDIA® BlueField®-2 DPU device.

1.1. Installation Notes

Refer to the [NVIDIA DOCA Installation Guide](#) for information on:

- ▶ Setting up NVIDIA® DOCA SDK on your NVIDIA® BlueField®-2 DPU
- ▶ Supported BlueField-2 platforms
- ▶ Supported host operating systems
- ▶ Package contents

1.2. Technical Support

Customers who purchased products directly from NVIDIA are invited to contact us on the [NVIDIA Support](#) website.

Chapter 2. Changes and New Features

- ▶ New libraries
 - ▶ Arg parser
 - ▶ Comm channel
 - ▶ DMA
 - ▶ RegEx
 - ▶ Core
- ▶ New applications
 - ▶ Allreduce – UCX based application
 - ▶ App Shield agent
 - ▶ Firewall - DOCA flow-based gRPC application
 - ▶ Secure channel
- ▶ New samples
 - ▶ AppShield
 - ▶ Comm channel
 - ▶ DMA
 - ▶ DPI
 - ▶ Flow
 - ▶ RegEx
 - ▶ Telemetry
- ▶ New services
 - ▶ Flow Inspector
 - ▶ HBN



Note: Compatible with DOCA 1.3 (BlueField software 3.9) and above.

- ▶ DOCA and CUDA development container
- ▶ DOCA and CUDA base runtime container

- ▶ DOCA and CUDA full runtime container
- ▶ New tools
 - ▶ apsh_config
- ▶ Libraries enhancements
 - ▶ Telemetry
 - ▶ Flow
 - ▶ HWS
 - ▶ Modes – VNF, Remote-VNF, Switch
 - ▶ Metadata
 - ▶ AppShield
 - ▶ Added beta-level support for AppShield library
- ▶ Application enhancements
 - ▶ Application Recognition – added support for connection tracking
 - ▶ URL Filter – added support for connection tracking
 - ▶ IPS – added support for connection tracking
 - ▶ DNS Filter
 - ▶ Added support for DOCA RegEx
 - ▶ Added support for NVIDIA® Converged Accelerator
 - ▶ East-west overlay encryption – added automation script
 - ▶ L4 OVS Firewall – added automation script
- ▶ Service enhancements
 - ▶ Telemetry
 - ▶ DOCA Development Container
 - ▶ DOCA Runtime Container
 - ▶ Application Recognition
 - ▶ URL Filter
 - ▶ IPS

Chapter 3. New Documentation

- ▶ DOCA Overview
- ▶ Programming Guides
 - ▶ Programming Guide Overview
 - ▶ Arg Parser Programming Guide
 - ▶ UCX Programming Guide
- ▶ Applications
 - ▶ Applications Overview
 - ▶ Allreduce
 - ▶ AppShield Agent
 - ▶ Firewall
 - ▶ Secure Channel
- ▶ Samples
 - ▶ AppShield
 - ▶ Comm Channel
 - ▶ DMA
 - ▶ DPI
 - ▶ Flow
 - ▶ RegEx
 - ▶ Telemetry Samples
- ▶ Tools
 - ▶ Tools Overview
- ▶ Miscellaneous
 - ▶ Troubleshooting
- ▶ Services
 - ▶ Services Overview
 - ▶ Flow Inspector

► HBN

Chapter 4. Known Issues

The following table lists the known issues and limitations for this release of DOCA SDK.

Reference	Description
3048250	Description: When configuring the DPU to operate in NIC Mode, the following parameters must be set to default (i.e., =0): HIDE_PORT2_PF, NVME_EMULATION_ENABLE, and VIRTIO_NET_EMULATION_ENABLE.
	Workaround: N/A
	Keyword: DPU operation mode
	Reported in version: 1.3
3049879	Description: When reloading (ifreload) an empty /etc/network/interfaces file, the previously created interfaces are not deleted.
	Workaround: To delete all previously created interfaces, at least one interface must be present in /etc/network/interfaces. The following configuration can be used as a safe "empty" file to delete all other virtual devices:
	<pre>auto p0 iface p0 auto p1 iface p1</pre>
	Keyword: HBN; unsupported NVUE commands
Reported in version: 1.3	
2821785	Description: Due to disabled backend foundation units, some commands show 500 INTERNAL SERVER ERROR/ 404 NOT FOUND. These commands are related to features or sub-systems which are not supported on HBN.
	Workaround: N/A
	Keyword: HBN; unsupported NVUE commands
	Reported in version: 1.3
2821785	Description: MAC addresses are not learned in the hardware but only in software. This may affect performance in pure L2 unicast traffic. This should not affect performance of IPv4/IPv6 traffic or L2 control traffic (i.e., STP, LLDP).
	Workaround: N/A
	Keyword: HBN
	Reported in version: 1.3
2828838	Description: NetworkManager and other services not directly related to HBN may display the following message in syslog:
	<pre>"netlink: read: too many netlink events. Need to resynchronize platform cache"</pre>

Reference	Description
	<p>The message has no functional impact and may be ignored.</p> <p>Workaround: N/A</p> <p>Keyword: HBN</p> <p>Reported in version: 1.3</p>
3042811	<p>Description: The DOCA API cannot be selected on the host (-d doca). RXPBench is not compiled against the DOCA API and therefore will fail to switch to using that API on the host. This option is available on the DPU.</p> <p>Workaround: After ensuring that all required DOCA libraries are installed, build RXPBench from the source on the host to enable this functionality.</p> <p>Keyword: RXPBench; host; tools</p> <p>Reported in version: 1.3</p>

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

© 2022 NVIDIA Corporation & affiliates. All rights reserved.