



# NVIDIA DOCA

## Release Notes

# Table of Contents

Chapter 1. New Features, Updates, and Enhancements.....	1
Chapter 2. New Documentation.....	3
Chapter 3. Installation Notes.....	4
3.1. DOCA Packages.....	4
3.2. Supported Operating System.....	5
3.3. Supported Kernel Versions.....	5
Chapter 4. Technical Support.....	6
Chapter 5. Known Issues.....	7

---

# Chapter 1. New Features, Updates, and Enhancements

- ▶ New DOCA services
  - ▶ FireFly
- ▶ New DOCA libraries
  - ▶ SHA
  - ▶ Rivermax (RMAX)
  - ▶ Compress
- ▶ DOCA Flow enhancements
  - ▶ Heterogeneous tables
  - ▶ Shared meters
  - ▶ Meter profile optimization
  - ▶ Ordered pipe list
- ▶ New DOCA reference applications
  - ▶ Switch
  - ▶ NAT
  - ▶ DMA Copy
  - ▶ File Compression
  - ▶ File Integrity
  - ▶ L2 Reflector
- ▶ New samples
  - ▶ Flow
    - ▶ Flow LPM
    - ▶ Flow Multi-actions
    - ▶ Flow RSS Meta
    - ▶ Flow Shared Counter

- ▶ Flow Shared Meter
- ▶ Rivermax (RMAX)
  - ▶ List Devices
  - ▶ Set CPU Affinity
  - ▶ Set Clock
  - ▶ Create Stream
  - ▶ Create Stream HDs
- ▶ New tools
  - ▶ Socket Relay
  - ▶ Comm channel admin tool
- ▶ FlexIO (BlueField-3 and emulation)
  - ▶ New DOCA FlexIO Driver
  - ▶ New tools: FlexIO Build, DPACC Compiler
- ▶ Service Enhancements
  - ▶ DOCA Telemetry Service – added counters and events
- ▶ Library enhancements
  - ▶ App Shield - privilege and environment variables
- ▶ Runtime Guide
  - ▶ TLS acceleration
  - ▶ Services Logger
- ▶ New DOCA Package for NVIDIA® ConnectX®

---

## Chapter 2. New Documentation

The following documents have been added to DOCA documentation:

- ▶ Services
  - ▶ Firefly
- ▶ Programming Guides
  - ▶ Compress
  - ▶ FlexIO
  - ▶ Rivermax (RMAX)
- ▶ Applications
  - ▶ DMA Copy
  - ▶ File Compression Application Guide
  - ▶ DOCA File Integrity
  - ▶ DOCA L2 Reflector
  - ▶ DOCA NAT
  - ▶ DOCA Switch
- ▶ Tools
  - ▶ Socket Relay
  - ▶ FlexIO Build
  - ▶ Comm Channel Admin Tool
  - ▶ DPACC Compiler
- ▶ Miscellaneous
  - ▶ TLS Offload
  - ▶ Services Fluent Logger



**Note:** Updates to existing documents are not listed here.

---

# Chapter 3. Installation Notes

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

- ▶ Setting up NVIDIA DOCA SDK on your BlueField DPU
- ▶ Supported BlueField platforms

## 3.1. DOCA Packages

Device	Component	Version	Description
Host	DOCA SDK	1.5.0	Software development kit package for developing host software
	DOCA Runtime	1.5.0	Runtime libraries required to run DOCA-based software applications on host
	DOCA Tools	1.5.0	Tools for developers and administrators on host
	Arm emulated (QEMU) development container	3.9.3	Linux-based BlueField Arm emulated container for developers
Target BlueField-2 DPU (Arm)	BlueField Software	3.9.3	BlueField image and firmware
	DOCA SDK	1.5.0	Software development kit packages for developing Arm software
	DOCA Runtime	1.5.0	Runtime libraries required to run DOCA-based software applications on Arm
	DOCA Tools	1.5.0	Tools for developers and administrators for Arm target

## 3.2. Supported Operating System

The operating system supported on the BlueField DPU is Ubuntu 20.04.

The following operating systems are supported on the host machine:

- ▶ Ubuntu 18.04/20.04/22.04
- ▶ CentOS/RHEL 7.6/8.0/8.2
- ▶ Rocky 8.6
- ▶ Debian 10.8

## 3.3. Supported Kernel Versions



**Note:** Only the following generic kernel versions are supported for DOCA local repo package for host installation (whether by SDKM or manually).

Host Operation System	Kernel Support	Arch Support
CentOS 7.6	4.14.0-115.el7a.aarch64	aarch64
	3.10.0-957.el7.x86_64	x86
CentOS 8.0	4.18.0-80.el8.x86_64	
CentOS 8.2	4.18.0-193.el8.x86_64	
RHEL 7.6	3.10.0-957.el7.x86_64	
RHEL 8.0	4.18.0-80.el8.x86_64	
RHEL 8.2	4.18.0-193.el8.x86_64	
Rocky 8.6	4.18.0-372.9.1.el8.x86_64	
Ubuntu 18.04	4.15.0-20-generic	
Ubuntu 20.04	5.4.0-26-generic	
Ubuntu 22.04	5.15.0-52-generic	
Debian 10.8	4.19.0-14-amd64	

---

## Chapter 4. Technical Support

Customers who purchased NVIDIA products directly from NVIDIA are invited to contact us through the following methods:

- ▶ E-mail: [enterprisesupport@nvidia.com](mailto:enterprisesupport@nvidia.com)
- ▶ Enterprise Support page: <https://www.nvidia.com/en-us/support/enterprise>

Customers who purchased NVIDIA M-1 Global Support Services, please see your contract for details regarding Technical Support.

Customers who purchased NVIDIA products through an NVIDIA-approved reseller should first seek assistance through their reseller.



---

# Chapter 5. Known Issues

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

Reference	Description
3168683	Description: If many interfaces are participating in EVPN/routing, it is possible for the routing process to run out of memory.
	Workaround: Have a maximum of 8 VF interfaces participating in routing/VXLAN.
	Keyword: HBN; routing; memory
	Reported in version: 1.5.0
3240785	Description: DTS might fail to connect to DPE if started after DTS is already running.
	Workaround: Start the DTS container only after starting DPE.
	Keyword: DTS; DPE; BlueMan
	Reported in version: 1.5.0
3250391	Description: <code>dpdk_queues_and_ports_init</code> fails when given 0 DPDK ports.
	Workaround: N/A
	Keyword: <code>doca_dpi_grpc</code> server; DPDK
	Reported in version: 1.5.0
3219539	Description: TC rules are programmed by OVS to map uplink and host representor ports to HBN service. These rules are ageable and can result in packets needing to get software forwarded periodically to refresh the rules.
	Workaround: The timeout value can be adjusted by changing the OVS parameter <code>other_config : max-idle</code> as documented <a href="#">here</a> . The shipped default value is 10000ms (10s).
	Keyword: HBN; SFC; aging
	Reported in version: 1.5.0
3184745	Description: The command <code>nv show interface &lt;intf&gt; acl</code> does not show correct information if there are multiple ACLs bound to the interface.
	Workaround: Use the command <code>nv show interface &lt;intf&gt;</code> to view the ACLs bound to an interface.
	Keyword: HBN; ACLs
	Reported in version: 1.5.0
3158934	Description: Deleting an NVUE user by removing their password file and restarting the <code>decrypt-user-add</code> service on the HBN container does not work.

Reference	Description
	<p>Workaround: Either respawn the container after deleting the file, or delete the password file corresponding to the user by running <code>userdel -r username</code>.</p> <p>Keyword: HBN; user deletion</p> <p>Reported in version: 1.5.0</p>
3191433	<p>Description: ECMP selection for the underlay path uses the ingress port and identifies uplink ports via round robin. This may not result in uniform spread of the traffic.</p> <p>Workaround: N/A</p> <p>Keyword: HBN; ECMP</p> <p>Reported in version: 1.5.0</p>
3185003	<p>Description: When a packet is encapsulated with a VXLAN header, it adds extra bytes which may cause the packet to exceed the MTU of link. Typically, the packet would be fragmented but its silently dropped and no fragmentation happens.</p> <p>Workaround: Make sure that the MTU on the uplink port is always 50 bytes more than host ports so that even after adding VXLAN headers, ingress packets do not exceed the MTU.</p> <p>Keyword: HBN; MTU; VXLAN</p> <p>Reported in version: 1.5.0</p>
3184905	<p>Description: On VXLAN encapsulation, the DF flag is not propagated to the outer header. Such a packet may be truncated when forwarded in the kernel, and it may be dropped when hardware offloaded.</p> <p>Workaround: Make sure that the MTU on the uplink port is always 50 bytes more than host ports so that even after adding VXLAN headers, ingress packets do not exceed the MTU.</p> <p>Keyword: HBN; VXLAN</p> <p>Reported in version: 1.5.0</p>
3188688	<p>Description: When stopping the container using the command <code>crictl stop</code> an error may be reported because the command uses a timeout of 0 which is not enough to stop all the processes in the HBN container.</p> <p>Workaround: Pass a timeout value when stopping the HBN container by running:  <pre>crictl stop --timeout 60 &lt;hbn-container&gt;</pre></p> <p>Keyword: HBN; timeout</p> <p>Reported in version: 1.5.0</p>
3129749	<p>Description: The same ACL rule cannot be applied in both the inbound and outbound direction on a port.</p> <p>Workaround: N/A</p> <p>Keyword: HBN; ACLs</p> <p>Reported in version: 1.5.0</p>
3126560	<p>Description: The system's time zone cannot be modified using NVUE in the HBN container.</p> <p>Workaround: The timezone can be manually changed by symlinking the <code>/etc/localtime</code> file to a binary time zone's identifier in the <code>/usr/share/zoneinfo</code> directory. For example:  <pre>sudo ln -sf /usr/share/zoneinfo/GMT /etc/localtime</pre></p>

Reference	Description
	Keyword: HBN; time zone; NVUE
	Reported in version: 1.5.0
3118204	Description: Auto-BGP functionality (where the ASN does not need to be configured but is dynamically inferred by the system based on the system's role as a leaf or spine device) is not supported on HBN.
	Workaround: If BGP is configured and used on HBN, the BGP ASN must be manually configured.
	Keyword: HBN; BGP
	Reported in version: 1.5.0
3233088	Description: Since checksum calculation is offloaded to the hardware (not done by the kernel), it is expected to see an incorrect checksum in the tcpdump for locally generated, outgoing packets. BGP keepalives and updates are some of the packets that show such incorrect checksum in tcpdump.
	Workaround: N/A
	Keyword: HBN; BGP
	Reported in version: 1.5.0
3240153	Description: DOCA kernel support only works on a non-default kernel.
	Workaround: N/A
	Keyword: Kernel
	Reported in version: 1.5.0
3239668	Description: The <code>12_reflector</code> reference application fails to start due to missing <code>libflexio.so</code> library.
	Workaround: Define <code>LD_LIBRARY_PATH</code> as specified in the <a href="#">NVIDIA DOCA Troubleshooting Guide</a> .
	Keyword: FlexIO; DOCA applications; <code>12_reflector</code>
	Reported in version: 1.5.0
3239630	Description: FlexIO sample <code>flexio_rpc</code> fails to compile.
	Workaround: Add the following at line 34 in the file <code>/opt/mellanox/doca/samples/flexio/flexio_rpc/meson.build</code> : <pre>source_file = meson.current_source_dir() + '/device/flexio_rpc_device.c'</pre>
	Keyword: FlexIO; DOCA samples; <code>flexio_rpc</code>
	Reported in version: 1.5.0
3217627	Description: The <code>doca_devinfo_rep_list_create</code> API returns success on the host instead of "Operation not supported".
	Workaround: N/A
	Keyword: DOCA core; InfiniBand
	Reported in version: 1.5.0
3048250	Description: When configuring the DPU to operate in NIC Mode, the following parameters must be set to default (i.e., =0): <code>HIDE_PORT2_PF</code> , <code>NVME_EMULATION_ENABLE</code> , and <code>VIRTIO_NET_EMULATION_ENABLE</code> .
	Workaround: N/A

Reference	Description
	Keyword: DPU operation mode
	Reported in version: 1.3.0
3049879	Description: When reloading (ifreload) an empty <code>/etc/network/interfaces</code> file, the previously created interfaces are not deleted.
	Workaround: To delete all previously created interfaces, at least one interface must be present in <code>/etc/network/interfaces</code> . 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: 1.3.0
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: 1.3.0
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: 1.3.0
2828838	Description: NetworkManager and other services not directly related to HBN may display the following message in syslog: "netlink: read: too many netlink events. Need to resynchronize platform cache" The message has no functional impact and may be ignored.
	Workaround: N/A
	Keyword: HBN
	Reported in version: 1.3.0

## 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.