



# NVIDIA DOCA

## Release Notes

# Table of Contents

Chapter 1. Introduction.....	1
Chapter 2. New Features, Updates, and Enhancements.....	2
Chapter 3. Installation Notes.....	3
3.1. Embedded DOCA Libraries.....	3
3.2. Embedded DOCA Drivers.....	3
3.3. DOCA Packages.....	6
3.4. Supported Operating System Distributions.....	7
Chapter 4. Technical Support.....	8
Chapter 5. Known Issues.....	9
Chapter 6. API Backward Compatibility.....	14
6.1. DOCA Core doca_buf.....	14

---

# Chapter 1. Introduction

DOCA 2.2.0 introduces enhancements for high-performance AI and bare metal clouds. With programmable congestion control and data path acceleration, DOCA SDK provides an extensive framework for developers.

---

# Chapter 2. New Features, Updates, and Enhancements



Important: This release only supports NVIDIA® BlueField®-2. For BlueField-3 please refer or upgrade to DOCA 2.2.1.



Important: Upgrading to this BSP version installs a new version of Ubuntu GRUB. This version of GRUB revokes the old UEFI secure boot certificates and install new ones. The new certificates will not validate older images and boot will fail. Therefore, to roll back to older software versions, users must disable UEFI secure boot.

- ▶ NIC mode enhancements for BlueField-3 – Higher performance and lower latency at scale using local DPU memory; performant RDMA with programmable congestion control (PCC)
- ▶ Added support for DPA with NIC mode
- ▶ Enabled PCC over BlueField-3 DPA (including in NIC mode) while allowing for up to 400Gb/s speeds
- ▶ Added support for a new DOCA profile, doca-ofed, for installing MLNX\_OFED with the DOCA package
- ▶ Added improvements and enhancements for DOCA Flow library
- ▶ Added beta-level support for OVS-DOCA, a highly optimized virtual switch for NVIDIA network services
- ▶ Redfish management – UEFI secure boot settings at scale, remotely, and securely
- ▶ DPU OS support – added Anolis OS 8.6 with kernel 5.10 support in BlueField-3 Arm
- ▶ Host OS support – added OS support: Debian 10.13, Allinux 3.2, OL 8.7, Ubuntu 22.04 (aarch64 host)



Important: No updates were made to the DOCA RegEx and DOCA DPI libraries in DOCA 2.2. Please refer to DOCA 2.5 for a note regarding future RegEx and DPI updates.

---

# 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. Embedded DOCA Libraries

Component	Version
doca-apps	2.2.0080-1
doca-grpc	2.2.0080-1
doca-libs	2.2.0080-1
ucx	1.15.0-1.2304052
gpunetio	2.2.0080-1

## 3.2. Embedded DOCA Drivers

Component	Version	Description
BlueField-3 firmware	<a href="#">32.38.1002</a>	Firmware is used to run user programs on the device which allow hardware to run
BlueField-2 firmware	<a href="#">24.38.1002</a>	Firmware is used to run user programs on the device which allow hardware to run
ATF	<a href="#">v2.2(release):4.2.0-80-g0ad5a0e</a>	Arm-trusted firmware is a reference implementation of secure world software for Arm architectures
UEFI	<a href="#">4.2.0-56-ge28fcb7</a>	UEFI is a specification that defines the architecture of the platform firmware used for booting and its interface for interaction with the operating system

Component	Version	Description
doca-base (MLNX_OFED)	<a href="#">23.07-0.5.0.0</a>	NVIDIA® MLNX_OFED is a single software stack that operates across all NVIDIA network adapter solutions
MFT	<a href="#">4.25.0-62</a>	NVIDIA® MFT is a set of firmware management and debug tools for NVIDIA devices
mlnx-dpdk	22.11.2307.2.0	Equivalent to DPDK upstream. The versioning of MLNX_DPDK indicates which upstream DPDK it is compatible with it (e.g., 22.11 is compatible with upstream DPDK 2022.11).
mlx-regex	1.2-ubuntu1	RegEx is a library that provides RegEx pattern matching to DOCA applications using the regular expression processor (RXP) or software-based engines when required
virtio-net-controller	1.6.14-1	Virtio-net-controller is a systemd service running on the DPU, with a user interface front-end to communicate with the background service
collectx-clxapi	1.13.2	A library which exposes the CollectX API, which allows any 3 <sup>rd</sup> party to easily use CollectX functionality in their own programs
libvma	9.8.31-1	The NVIDIA® VMA library accelerates latency-sensitive and throughput-demanding TCP and UDP socket-based applications by offloading traffic from the user-space directly to the NIC, without going through the kernel and the standard IP stack (kernel-bypass)
libxlio	3.10.5-1.2307050	The NVIDIA® XLIO software library boosts the performance of TCP/IP applications based on NGINX (CDN, DoH, etc.) and storage solutions as part of the SPDK
dpcp	1.1.40-1.2307050	DPCP provides a unified flexible interface for programming IB devices using DevX

Component	Version	Description
mlnx-snap	3.7.4-2	BlueField SNAP for NVMe and virtio-blk enables hardware-accelerated virtualization of local storage
mlnx-libsnap	1.5.4-7	Libsnap is a common library designed to assist common tasks for applications wishing to interact with emulated hardware over BlueField DPUs and take the most advantage from hardware capabilities
spdk	23.01-11	SPDK provides a set of tools and libraries for writing high performance, scalable, user-mode storage applications
flexio	23.7.1694	FlexIO SDK exposes an API for managing the device and executing native code over the DPA processor
dpacc	1.5.0	DPACC is a high-level compiler for the DPA processor which compiles code targeted for the data-path accelerator (DPA) processor into a device executable and generates a DPA program
rxp-compiler	23.07.1	NVIDIA® RXP® is a processor developed to efficiently process data to detect matches for a set of user-defined string and regular expression (RegEx) based rules and is used to compile regular expressions into a format that can be executed by the RXP
rxpbench	2.2.0080-1	RXPBench is a tool that allows for the performance comparison between the NVIDIA® RXP® hardware RegEx acceleration engine found in the BlueField DPU and the Intel® Hyperscan software library. It provides a comprehensive set of options and can facilitate ingress of data from live network ports or previously recorded PCAP files.

Component	Version	Description
Rivermax	1.31	NVIDIA® Rivermax® is an optimized networking SDK for media and data streaming applications
RShim	2.0.9.gb35835f_amd64	The user-space driver to access the BlueField SoC via the RShim interface, providing ways to push boot stream, debug the target, or login via the virtual console or network interface

### 3.3. DOCA Packages

Device	Component	Version	Description
Host	DOCA SDK	2.2.0	Software development kit package for developing host software
	DOCA Runtime	2.2.0	Runtime libraries required to run DOCA-based software applications on host
	DOCA Tools	2.2.0	Tools for developers and administrators on host
	DOCA Extra	2.2.0	Contains helper scripts (doca-info, doca-kernel-support)
	DOCA OFED	2.2.0	Single software stack which operates across all NVIDIA network adapter solutions
	Arm emulated (QEMU) development container	4.2.0	Linux-based BlueField Arm emulated container for developers
Target BlueField DPU (Arm)	BlueField BSP	4.2.0	BlueField image and firmware
	DOCA SDK	2.2.0	Software development kit packages for developing Arm software
	DOCA Runtime	2.2.0	Runtime libraries required to run DOCA-



Device	Component	Version	Description
	DOCA Tools	2.2.0	based software applications on Arm Tools for developers and administrators for Arm target

## 3.4. Supported Operating System Distributions

The default operating system of the BlueField DPU (Arm) is Ubuntu 22.04.

The supported operating systems on the host machine are the following:



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

OS	Kernel	x86	aarch64
Alinux 3.2	5.10	#	
BCLinux 21.10 SP2	4.19.90	#	#
CTyunOS 2	4.19.90	#	#
Debian 10.13	5.10.135	#	
	5.4.210		
Debian 10.8	4.19.0	#	
Kylin 10 SP2	4.19.90	#	#
openEuler 20.03 SP3	4.19.90	#	#
openEuler 22.03	5.10.0	#	#
Oracle Linux 8.7	5.10/5.15	#	
RHEL/CentOS 7.6	3.10	#	
	4.14		#
RHEL/CentOS 8.0	4.18	#	
RHEL/CentOS 8.2	4.18	#	
RHEL/Rocky 8.6	4.18	#	
RHEL/Rocky 9.1	5.14.0	#	#
Ubuntu 18.04	4.15	#	
Ubuntu 20.04	5.4	#	
Ubuntu 22.04	5.15	#	#
UOS 20 1040d	4.19.0	#	#
Windows	DOCA support for Windows is on DOCA's roadmap. For Windows driver support, refer to <a href="#">WinOF-2 Release Notes</a> .		

---

# 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
3566042	Description: Virtio hotplug is not supported in GPU-HOST mode on the NVIDIA Converged Accelerator.
	Workaround: N/A
	Keyword: Virtio; Converged Accelerator
	Reported in version: 2.2.0
3546474	Description: PXE boot over ConnectX interface might not work due to an invalid MAC address in the UEFI boot entry.
	Workaround: On the DPU, create <code>/etc/bf.cfg</code> file with the relevant PXE boot entries, then run the command <code>bfcfg</code> .
	Keyword: PXE; boot; MAC
	Reported in version: 2.2.0
3393316	Description: When LSO is enabled, if the header and data appear in the same fragment, the following warning is given from tcpdump: <code>truncated-ip - 9 bytes missing</code>
	Workaround: N/A
	Keyword: Virtio-net; large send offload
	Reported in version: 2.2.0
3549785	Description: NVMe and <code>mlx5_core</code> drivers fail during BFB installation. As a result, Anolis OS cannot be installed on the SSD and the <code>mlxfwreset</code> command does not work during Anolis BFB installation.
	Workaround: N/A
	Keyword: Linux; NVMe; BFB installation
	Reported in version: 2.2.0
3561723	Description: Running <code>mlxfwreset sync 1</code> on NVIDIA Converged Accelerators may be reported as supported although it is not. Executing the reset will fail.
	Workaround: N/A
	Keyword: <code>mlxfwreset</code>
	Reported in version: 2.2.0

Reference	Description
3306489	<p>Description: After rebooting a BlueField-3 DPU running Rocky Linux 8.6 BFB, the kernel log shows the following error:</p> <pre>[ 3.787135] mlxbf_gige MLNXBF17:00: Error getting PHY irq. Use polling instead</pre> <p>This message indicates that the Ethernet driver will function normally in all aspects, except that PHY polling is enabled.</p> <p>Workaround: N/A</p> <p>Keyword: Linux; PHY; kernel</p> <p>Reported in version: 2.2.0</p>
3306489	<p>Description: When performing longevity tests (e.g., mlxfwreset, DPU reboot, burning of new BFBs), a host running an Intel CPU may observe errors related to "CPU 0: Machine Check Exception".</p> <p>Workaround: Add <code>intel_idle.max_cstate=1</code> entry to the kernel command line.</p> <p>Keyword: Longevity; mlxfwreset; DPU reboot</p> <p>Reported in version: 2.2.0</p>
3529297	<p>Description: Enhanced NIC mode is not supported on BlueField-2 DPUs.</p> <p>Workaround: N/A</p> <p>Keyword: Operation; mode</p> <p>Reported in version: 2.2.0</p>
3538486	<p>Description: When removing LAG configuration from the DPU, a kernel warning for <code>uverbs_destroy_ufile_hw</code> is observed if <code>virtio-net-controller</code> is still running.</p> <p>Workaround: Stop <code>virtio-net-controller</code> service before cleaning up bond configuration.</p> <p>Keyword: Virtio-net; LAG</p> <p>Reported in version: 2.2.0</p>
3527302	<p>Description: Failure occurs on <code>doca_mmap_start()</code> if the memory range is from <code>dmabuf</code> (i.e., if <code>mmap</code> is created with <code>doca_mmap_set_dmabuf_memrange()</code> call).</p> <p>Workaround: N/A</p> <p>Keyword: Memory</p> <p>Reported in version: 2.2.0</p>
3541010	<p>Description: In case of an asynchronous wait, submitting a <code>doca_sync_event_job_wait</code> job is limited to a Sync Event with a value in the range [0, 254] and is limited to a wait threshold in the range [0,254]. Other scenarios result in anomalous behavior.</p> <p>Workaround: N/A</p> <p>Keyword: Sync-event; kernel</p> <p>Reported in version: 2.2.0</p>
3511313	<p>Description: On BlueField-3, the MAC addresses of Arm ports (p0 and p1) do not match the value on DPU sticker.</p> <p>Workaround: N/A</p> <p>Keyword: Port; MAC address</p>

Reference	Description
	Reported in version: 2.2.0
3533508	Description: OVS-dpdk is not supported if grub is used to allocate hugepages. Workaround: N/A Keyword: Hugepages; OVS-DPDK Reported in version: 2.2.0
3533850	Description: PCC is not supported when operating in DPU mode. Workaround: N/A Keyword: PCC Reported in version: 2.2.0
3534219	Description: On BlueField-3 devices, from DOCA 2.2.0 to 32.37.1306 (or lower), the host crashes when executing partial Arm reset (e.g., Arm reboot; BFB push; mlxfwreset). Workaround: Before downgrading the firmware: 1. Run: <pre>echo 0 &gt; /sys/bus/platform/drivers/mlxbf-bootctl/large_icm</pre> 2. Reboot Arm. Keyword: PCC; hang Reported in version: 2.2.0
3530300	Description: DOCA_PCC application may be terminated due to a false hang monitor alarm when running traffic. Workaround: N/A Keyword: PCC; hang Reported in version: 2.2.0
N/A	Description: The <a href="#">NVIDIA DOCA East-West Overlay Encryption Application</a> (and the underlying DPU OS Kernel driver IPsec functionality) is not supported. User space DOCA IPsec is not impacted. Workaround: N/A Keyword: IPsec Reported in version: 2.2.0
3382740	Description: Fragmented packets are not supported in <a href="#">Application Recognition</a> , <a href="#">Intrusion Prevention</a> , and <a href="#">URL Filtering</a> reference applications. Workaround: N/A Keyword: Fragmented packets; DOCA applications Reported in version: 2.2.0
3444073	Description: mlxfwreset is not supported in this release. Workaround: Power cycle the host. Keyword: mlxfwreset; support Reported in version: 2.0.2

Reference	Description
3448841	Description: While running CentOS 8.2, switchdev Ethernet DPU runs in "shared" RDMA net namespace mode instead of "exclusive".
	Workaround: Use <code>ib_core</code> module parameter <code>netns_mode=0</code> . For example: <pre>echo "options ib_core netns_mode=0" &gt;&gt; /etc/modprobe.d/mlnx-bf.conf</pre>
	Keyword: RDMA; isolation; Net NS
	Reported in version: 2.0.2
3365363	Description: On BlueField-3, when booting virtio-net emulation device using a GRUB2 bootloader, the bootloader may attempt to close and re-open the virtio-net device. This can result in unexpected behavior and possible system failure to boot.
	Workaround: N/A
	Keyword: BlueField-3; virtio-net; UEFI
	Reported in version: 2.0.2
3232444	Description: After live migration of virtio-net devices using the VFE driver, the <code>max_queues_size</code> output from the <code>virtnet list</code> may be wrong. This does not affect the actual value.
	Workaround: N/A
	Keyword: Virtio-net; live migration
	Reported in version: 2.0.2
3441287	Description: Failure occurs when attempting to raise static LAG with <code>ifenslave_2.10ubuntu3</code> package.
	Workaround: Use <code>ifenslave_2.9ubuntu1</code> .
	Keyword: <code>ifenslave</code> ; bonding
	Reported in version: 2.0.2
3373849	Description: Different OVS-based packages can include their own <code>systemd</code> services which prevents <code>/sbin/mlnx_bf_configure</code> from identifying the right one.
	Workaround: Use a specific service name in <code>/sbin/mlnx_bf_configure</code> .
	Keyword: OVS; <code>systemd</code>
	Reported in version: 2.0.2
2706803	Description: When an NVMe controller, SoC management controller, and DMA controller are configured, the maximum number of VFs is limited to 124.
	Workaround: N/A
	Keyword: VF; limitation
	Reported in version: 2.0.2
3380586	Description: Public key acceleration is not enabled on OpenEuler BFB due to missing configurations in the <code>openssl.cnf</code> file.
	Workaround: N/A
	Keyword: PKA; OpenSSL
	Reported in version: 2.0.2
3273435	Description: Changing the mode of operation between NIC and DPU modes results in different capabilities for the host driver which might cause unexpected behavior.
	Workaround: Reload the host driver or reboot the host.

Reference	Description
	Keyword: Modes of operation; driver
	Reported in version: 2.0.2
3438222	Description: On BlueField DPU running Rocky, openEuler or Centos8.2 with default huge page size not equal to 2M, rxpbench fails to initialize as no mounted <code>hugetlbfs</code> is found for the 2M size.
	Workaround: N/A
	Keyword: rxpbench
	Reported in version: 2.0.2
3377199	Description: After installing OpenEuler 20.03sp1 BFB, the 2nd port may raise configured with legacy mode.
	Workaround: Reboot the DPU.
	Keyword: OpenEuler; legacy
	Reported in version: 2.0.2
3362822	Description: Running the gRPC firewall application in interactive mode and trying to add and remove the same entry leads to a failure on the server side which severs the connection to the server.
	Workaround: N/A
	Keyword: gRPC; firewall; interactive; connection
	Reported in version: 2.0.2
3264749	Description: In Rocky and CentOS 8.2 inbox-kernel BFBs, RegEx requires the following extra huge page configuration for it to function properly: <pre>sudo hugeadm --pool-pages-min DEFAULT:2048M sudo systemctl start mlx-regex.service systemctl status mlx-regex.service</pre> If these commands have executed successfully you should see <code>active (running)</code> in the last line of the output.
	Workaround: N/A
	Keyword: RegEx; hugepages
	Reported in version: 1.5.1
3240153	Description: DOCA kernel support only works on a non-default kernel.
	Workaround: N/A
	Keyword: Kernel
	Reported in version: 1.5.0
3217627	Description: The <code>doca_devinfo_rep_list_create</code> API returns success on the host instead of <code>Operation not supported</code> .
	Workaround: N/A
	Keyword: DOCA core; InfiniBand
	Reported in version: 1.5.0

---

# Chapter 6. API Backward Compatibility

This section lists changes to the DOCA SDK which impacts backward compatibility.

## 6.1. DOCA Core `doca_buf`

Up to DOCA 2.0.2, the data length of the buffer is ignored when using the buffer as an output parameter, and the new data was written over the data that was there beforehand. From now on, new data is appended after existing data (if any) while updating the data length accordingly.

Because of this change, it is recommended that a destination buffer is allocated without a data section (data length 0), for ease of use.

In cases where the data length is 0 in a destination buffer, this change would go unnoticed (as appending the data and writing to the data section has the same result).

Reusing buffers requires resetting the data length when wishing to write to the same data address (instead of appending the data), overwriting the existing data. A new function, `doca_buf_reset_data_len()`, has been added specifically for this need.



## 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.