



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 1. Introduction

NVIDIA DOCA™ 2.2.1 is a special GA software release meant to support and stabilize NVIDIA® BlueField®-3 DPUs only. For a DOCA version that supports BlueField-2 DPUs and ConnectX family adapter cards, please refer to DOCA 2.2.0.

Chapter 2. New Features, Updates, and Enhancements

- ▶ Added support for signing and authenticating user DPA applications for BlueField-3, allowing only signed applications to load on DPA
- ▶ Added Redfish support for configuring all UEFI secure boot settings (disable, enable, enroll user keys, etc.) at scale, remotely, and securely
- ▶ Added alpha support for BlueMan



WARNING: The webserver for BlueMan Alpha uses HTTP connection. Users are advised to reconfigure the webserver to HTTPS.

- ▶ Added support for the DOCA PCC library to support tracing on the device side of the application code using the call `doca_pcc_dev_trace_5(...)`. The common use case for this call is to trace a specific flow using its flow tag, which can be acquired by the `doca_pcc_dev_get_flowtag()` call. This trace call is designed to be faster than the normal `doca_pcc_dev_printf()` call, and is added as a debugging tool that does not highly affect performance.
- ▶ Individual priorities in `mlxconfig`
- ▶ Optimized DPA RTOS performance



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	32.38.3056	Firmware is used to run user programs on the device which allow hardware to run
ATF	v2.2(release): 4.2.2-3-ged2a420	Arm-trusted firmware is a reference implementation of secure world software for Arm architectures
UEFI	4.2.2-2-g1174066	UEFI is a specification that defines the architecture of the platform firmware used for booting and its interface for interaction with the operating system
doca-base (MLNX_OFED)	23.07-0.5.0.0	NVIDIA® MLNX_OFED is a single software stack that

Component	Version	Description
MFT	4.25.1-11	operates across all NVIDIA network adapter solutions 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	1.6.27-1	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 rd 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
mlnx-snap	3.7.4-2	BlueField SNAP for NVMe and virtio-blk enables hardware-

Component	Version	Description
mlnx-libsnap	1.5.4-9	accelerated virtualization of local storage 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	23.9.1774	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.
Rivermax	1.31	NVIDIA® Rivermax® is an optimized networking SDK for media and data streaming applications

Component	Version	Description
RShim	2.0.9-0.gb35835f	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	2.2.1	Software development kit package for developing host software
	DOCA Runtime	2.2.1	Runtime libraries required to run DOCA-based software applications on host
	DOCA Tools	2.2.1	Tools for developers and administrators on host
	DOCA Extra	2.2.1	Contains helper scripts (doca-info, doca-kernel-support)
	DOCA OFED	2.2.1	Single software stack which operates across all NVIDIA network adapter solutions
	Arm emulated (QEMU) development container	4.2.2	Linux-based BlueField Arm emulated container for developers
Target BlueField DPU (Arm)	BlueField BSP	4.2.2	BlueField image and firmware
	DOCA SDK	2.2.1	Software development kit packages for developing Arm software
	DOCA Runtime	2.2.1	Runtime libraries required to run DOCA-based software applications on Arm
	DOCA Tools	2.2.1	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	#	
Debian 10.13	5.10.135	#	
	5.4.210		
Debian 10.8	4.19.0	#	
Oracle Linux 8.7	5.10/5.15	#	
RHEL/CentOS 7.6	3.10	#	
RHEL/CentOS 8.2	4.18	#	
Ubuntu 18.04	4.15	#	
Ubuntu 20.04	5.4	#	
Ubuntu 22.04	5.15	#	
Windows	DOCA support for Windows is on DOCA's roadmap. For Windows driver support, refer to WinOF-2 Release Notes .		

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
- ▶ 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
3603146	Description: Running <code>mlxfwreset</code> on BlueField-3 may cause the external host to crash when the RShim driver is running on that host.
	Workaround: Stop the RShim driver on the external host using <code>systemctl stop rshim</code> before performing <code>mlxfwreset</code> .
	Keyword: RShim; mlxfwreset
	Reported in version: 2.2.1
3594836	Description: When enabling Flex IO SDK tracer at high rates, a slow-down in processing may occur and/or some traces may be lost.
	Workaround: Keep tracing limited to ~1M traces per second to avoid a significant processing slow-down. Use tracer for debug purposes and consider disabling it by default.
	Keyword: Tracer FlexIO
	Reported in version: 2.2.1
3556795	Description: The first uplink representor interface may not be renamed to p0 from ethX.
	Workaround: Run the command <code>udevadm trigger -s net -c add -v</code> .
	Keyword: Representors
	Reported in version: 2.2.1
3592080	Description: When using UEK8 on the host in DPU mode, creating a VF on the host consumes about 100MB memory on the DPU.
	Workaround: N/A
	Keyword: UEK; VF
	Reported in version: 2.2.1
3568924	Description: <code>doca_rdma</code> does not support cross-subnet routing.
	Workaround: N/A
	Keyword: RDMA
	Reported in version: 2.2.1
3566042	Description: Virtio hotplug is not supported in GPU-HOST mode on the NVIDIA Converged Accelerator.
	Workaround: N/A

Reference	Description
	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
3306489	Description: After rebooting a BlueField-3 DPU running Rocky Linux 8.6 BFB, the kernel log shows the following error: <code>[3.787135] mlxbf_gige MLNXBF17:00: Error getting PHY irq. Use polling instead</code>
	This message indicates that the Ethernet driver will function normally in all aspects, except that PHY polling is enabled.
	Workaround: N/A
	Keyword: Linux; PHY; kernel
	Reported in version: 2.2.0
3306489	Description: When performing longevity tests (e.g., <code>mlxfwreset</code> , DPU reboot, burning of new BFBs), a host running an Intel CPU may observe errors related to "CPU 0: Machine Check Exception".
	Workaround: Add <code>intel_idle.max_cstate=1</code> entry to the kernel command line.
	Keyword: Longevity; <code>mlxfwreset</code> ; DPU reboot
	Reported in version: 2.2.0
3529297	Description: Enhanced NIC mode is not supported on BlueField-2 DPUs.

Reference	Description
	Workaround: N/A
	Keyword: Operation; mode
	Reported in version: 2.2.0
3538486	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.
	Workaround: Stop <code>virtio-net-controller</code> service before cleaning up bond configuration.
	Keyword: Virtio-net; LAG
	Reported in version: 2.2.0
3527302	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).
	Workaround: N/A
	Keyword: Memory
	Reported in version: 2.2.0
3541010	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 <code>[0, 254]</code> and is limited to a wait threshold in the range <code>[0,254]</code> . Other scenarios result in anomalous behavior.
	Workaround: N/A
	Keyword: Sync-event; kernel
	Reported in version: 2.2.0
3511313	Description: On BlueField-3, the MAC addresses of Arm ports (p0 and p1) do not match the value on DPU sticker.
	Workaround: N/A
	Keyword: Port; MAC address
	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; <code>mlxfwreset</code>).
	Workaround: Before downgrading the firmware:
	1. Run:
	<code>echo 0 > /sys/bus/platform/drivers/mlxbf-bootctl/large_icm</code>

Reference	Description
	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 NVIDIA DOCA East-West Overlay Encryption Application (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 Application Recognition , Intrusion Prevention , and URL Filtering reference applications.
	Workaround: N/A
	Keyword: Fragmented packets; DOCA applications
	Reported in version: 2.2.0
3444073	Description: <code>mlxfwreset</code> is not supported in this release.
	Workaround: Power cycle the host.
	Keyword: <code>mlxfwreset</code> ; support
	Reported in version: 2.0.2
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" >> /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

Reference	Description
	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.
	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, <code>rxpbench</code> fails to initialize as no mounted <code>hugetlbfs</code> is found for the 2M size.
	Workaround: N/A
	Keyword: <code>rxpbench</code>
	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.

Reference	Description
	Workaround: N/A
	Keyword: gRPC; firewall; interactive; connection
	Reported in version: 2.0.2
3264749	<p>Description: In Rocky and CentOS 8.2 inbox-kernel BFBs, RegEx requires the following extra huge page configuration for it to function properly:</p> <pre>sudo hugeadm --pool-pages-min DEFAULT:2048M sudo systemctl start mlx-regex.service systemctl status mlx-regex.service</pre> <p>If these commands have executed successfully you should see <code>active (running)</code> in the last line of the output.</p>
	Workaround: N/A
	Keyword: RegEx; hugepages
	Reported in version: 1.5.1
3240153	<p>Description: DOCA kernel support only works on a non-default kernel.</p>
	Workaround: N/A
	Keyword: Kernel
	Reported in version: 1.5.0
3217627	<p>Description: The <code>doca_devinfo_rep_list_create</code> API returns success on the host instead of <code>Operation not supported</code>.</p>
	Workaround: N/A
	Keyword: DOCA core; InfiniBand
	Reported in version: 1.5.0

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

© 2023 NVIDIA Corporation & affiliates. All rights reserved.