



NVIDIA DOCA Release Notes

Table of contents

General Support	6
Changes and New Features	26
Bug Fixes in This Version	28
Known Issues	32

NVIDIA DOCA SDK release notes containing information on new features, software interoperability, and known issues.

1. Introduction

DOCA 2.8.0 introduces NVIDIA® BlueField® networking platform (DPU or SuperNIC) enhancements for high-performance and secure AI bare-metal cloud and DOCA-Host updates for supported BlueField and NVIDIA® ConnectX® devices. With programmable congestion control (PCC) and data-path acceleration (DPA). DOCA SDK provides an extensive framework for developers.

The DOCA release notes contain the following subpages:

- [General Support](#)
- [Changes and New Features](#)
- [Bug Fixes in This Version](#)
- [Known Issues](#)

2. Installation Notes

Note

BlueField-3 devices are not supported with MLNX_OFED as the host driver and are required to use DOCA-Host.

Warning

BlueField DPUs with the following SKUs require an 8-pin ATX power supply cable connection when powering up. Without this connection to the power supply cable, the device will not complete the power-on procedure and will not function properly.

- *B3220 DPUs – 900-9D3B6-00CV-AA0 and 900-9D3B6-00SV-AA0

- *B3240 DPUs – 900-9D3B6-00CN-AB0 and 900-9D3B6-00SN-AB0
- *B3210 DPUs – 900-9D3B6-00CC-AA0 and 900-9D3B6-00SC-AA0
- *B3210E DPUs – 900-9D3B6-00CC-EA0 and 900-9D3B6-00SC-EA0

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

- Setting up DOCA SDK on your BlueField networking platform or SmartNIC
- Supported BlueField platforms

(i) Note

By default, installing DOCA profiles with standard Linux tools (yum, apt) installs both `doca-runtime` and `doca-devel` (previously `doca-sdk`).

- `doca-runtime` includes all the components, libs, drivers, and tools used in the production environment by the DOCA admin
- `doca-devel` includes all the components, libs, drivers, and tools used for development, including reference applications, compilers, etc.

Starting with DOCA 2.8.0, the default installation of BlueField-Bundle and DOCA-Host profiles will only include DOCA runtime.

`doca-devel` can be installed manually as needed.

3. Supported Device Speeds

Uplink/Adapter Card	Driver Name	Uplink Speed
BlueField-2	mlx5	<ul style="list-style-type: none"> InfiniBand: SDR, FDR, EDR, HDR Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE ¹, 100GbE ¹
BlueField		<ul style="list-style-type: none"> InfiniBand: SDR, QDR, FDR, FDR10, EDR Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 100GbE
ConnectX-7		<ul style="list-style-type: none"> InfiniBand: EDR, HDR100, HDR, NDR200, NDR Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE ¹, 100GbE ¹, 200GbE ², 400GbE
ConnectX-6 Lx		<ul style="list-style-type: none"> Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE ¹
ConnectX-6 Dx		<ul style="list-style-type: none"> Ethernet: 10GbE, 25GbE, 40GbE, 50GbE ¹, 100GbE ¹, 200GbE ¹
ConnectX-6		<ul style="list-style-type: none"> InfiniBand: SDR, FDR, EDR, HDR Ethernet: 10GbE, 25GbE, 40GbE, 50GbE ¹, 100GbE ¹, 200GbE ¹
ConnectX-5/ConnectX-5 Ex		<ul style="list-style-type: none"> InfiniBand: SDR, QDR, FDR, FDR10, EDR Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 100GbE
ConnectX-4 Lx		<ul style="list-style-type: none"> Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE
ConnectX-4		<ul style="list-style-type: none"> InfiniBand: SDR, QDR, FDR, FDR10, EDR Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 56GbE ³, 100GbE

1. Speed that supports both NRZ and PAM4 modes in Force mode and Auto-Negotiation mode. □ □ □ □ □ □ □ □ □ □ □ □

2. Speed that supports PAM4 mode only. □

3. 56GbE is an NVIDIA proprietary link speed and can be achieved while connecting an NVIDIA adapter card to NVIDIA SX10XX switch series or when connecting an NVIDIA adapter card to another NVIDIA adapter card. [□](#)

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.

Info

For questions, comments, and feedback, please contact us at DOCA-Feedback@exchange.nvidia.com.

General Support

1. Embedded DOCA Firmware Components

Component	Version	Description
ATF	v2.2(release):4.8.0-41-gf0ff3a4	Arm-trusted firmware is a reference implementation of secure world software for Arm architectures
UEFI	4.8.0-36-gf01f42f	UEFI is a specification that defines the architecture of the platform firmware used for booting and its interface for interaction with the operating system
BlueField-3 NIC firmware	32.42.1000	Firmware is used to run user programs on the BlueField-3 which allow hardware to run
BlueField-2 NIC firmware	24.42.1000	Firmware is used to run user programs on the BlueField-2 which allow hardware to run
BMC firmware	24.07	BlueField BMC firmware
BlueField-3 eROT (Glacier)	00.02.0182.0000	BlueField-3 eROT firmware
BlueField-2 eROT (CEC)	cec_ota_BMGP-04.0f	BlueField-2 eROT firmware

2. Supported NIC Firmware Versions

Note

DOCA 2.9.0 will be the last DOCA release to support ConnectX-4. DOCA 2.9.0 will be an LTS version and will be supported for 3 years for bug fixes and CVE updates.

Adapter Card	Bundled Firmware Version
BlueField-2	24.42.1000
ConnectX-7	28.42.1000
ConnectX-6 Lx	26.42.1000
ConnectX-6 Dx	22.42.1000
ConnectX-6	20.42.1000
ConnectX-5/ConnectX-5 Ex	16.35.4030
BlueField	18.33.1048
ConnectX-4 Lx	14.32.1010
ConnectX-4	12.28.2006

To obtain the official firmware versions, refer to the NVIDIA [Firmware Download](#) page.

3. Embedded DOCA Drivers

Component	Version	Description	Licens
clusterkit	1.14.462-1.2407052	A multifaceted node assessment tool for high-performance clusters	BSD
collectx-clxapi	1.18.2-17111037	A library which exposes the CollectX API, which allows any 3 rd party to easily use CollectX functionality in their own programs	Proprie
dpacc	1.8.0	DPACC is a high-level compiler for the DPA processor which compiles code targeted for the data-path accelerator (DPA)	Proprie

Component	Version	Description	Licens
		processor into a device executable and generates a DPA program	
dpcp	1.1.49-1.2407052	DPCP provides a unified flexible interface for programming IB devices using DevX	Proprie
flexio	24.04.2148-0	FlexIO SDK exposes an API for managing the device and executing native code over the DPA processor	Proprie
fwctl	24.07-OFED.24.07.0.5.1.1	Subsystem designed to standardize the secure firmware interface for userspace, focusing on debugging, configuration, and provisioning	GPLv2
hcoll	4.8.3228-1.2407052	HCOLL contains support for building runtime configurable hierarchical collectives	Proprie
ibarr	0.1.3-1.2407052	ip2gid address resolution and gid2lid path record resolution	GPL-2. with Li syscall or BSD Clause
ibdump	6.0.0-1.2407052	Dump of InfiniBand traffic;	BSD2+

Component	Version	Description	Licens
		diagnostic tool	
ibsim	0.12-1.2407052	Open-source InfiniBand fabric simulator	GPLv2 BSD
ibutils	2.1.1	ibdiagnet scans the fabric using directed route packets and extracts all the available information regarding its connectivity and devices.	Proprie
ibutils2	2.1.1-0.21800.MLNX20240801.ga4352587.2407052	Utilities for InfiniBand	Proprie
iser	24.07-OFED.24.07.0.5.2.1	Storage related drivers	GPLv2
isert	24.07-OFED.24.07.0.5.2.1	Storage related drivers	GPLv2
kernel-mft	4.29.0-127	Kernel part of MFT tools (for firmware burning, etc.)	Dual BSD/G
knem	1.1.4.90mlnx3-OFED.23.10.0.2.1.1	Open-source kernel module that enables high-perf intra-node MPI communication	BSD an GPLv2
libvma	9.8.60-1	The NVIDIA® Messaging Accelerator (VMA) library accelerates latency-sensitive and throughput-demanding TCP and UDP socket-based applications by offloading traffic from the	GPLv2 BSD

Component	Version	Description	Licens
		user-space directly to the NIC, without going through the kernel and the standard IP stack (kernel-bypass)	
libxlio	3.31.2-1	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	GPLv2 BSD
mft	4.29.0-131	NVIDIA® MFT is a set of firmware management and debug tools for NVIDIA devices	Proprietary
mlnx-dpdk	22.11.0-2404	Equivalent to DPDK upstream. The versioning of MLNX_DPDK indicates which upstream DPDK it is compatible with (e.g., 22.11 is compatible with upstream DPDK 2022.11).	BSD, LGPLv3, and GPL
mlnx-en	24.07-0.5.2.0.ge08362d	Kernel drivers part for Ethernet-only package	GPLv2
mlnx-ethtool	24.07-0.5.2.0.ge08362d	Ethtool with optional MLNX adaptation	GPL

Component	Version	Description	Licens
mlnx- iproute2	6.9.0-1.2407052	IPRoute with optional MLNX adaptation	GPL
mlnx- libsnap	1.6.0-1	Libsnap is a common library designed to assist common tasks for applications wishing to interact with emulated hardware over BlueField and take the most advantage from hardware capabilities	Proprie
mlnx- nfsrdma	24.07-OFED.24.07.0.5.2.1	Storage related driver for NFS over RDMA	GPLv2
mlnx-nvme	24.07-OFED.24.07.0.5.2.1	Storage related driver for NVMe	GPLv2
mlnx- ofa_kernel	24.07-OFED.24.07.0.5.2.1	Kernel drivers for Ethernet InfiniBand together	GPLv2
mlnx-snap	3.8.0-3	BlueField SNAP for NVMe and virtio-blk enables hardware-accelerated virtualization of local storage	Proprie
mlnx-tools	24.07-0.2407052	Tools for loading modules, configurations, scripts, etc.	GPLv2 BSD
mlx-regex	1.2-ubuntu1	RegEx is a library that provides RegEx pattern	Proprie

Component	Version	Description	Licens
		matching to DOCA applications using the regular expression processor (RXP) or software-based engines when required	
mlx-steering-dump	1.0.0-0.2407052	Hardware/software steering dump parsing tools	GPLv2
mpitests	3.2.24-2ffc2d6.2407052	Test suite for benchmarking the MPI	BSD
mstflint	4.26.0-1	User space part of our MFT tools	GPL/B:
multiperf	3.0-3.0.2407052	Linux tool for perf testing	BSD 3- Clause v2 or l
ofed-scripts	24.07-OFED.24.07.0.5.2	Scripts used to build OFED	GPL/B:
openmpi	4.1.7a1-1.2407052	MPI implementation (for RDMA/RoCE) with some improvements done by the HPC team	BSD
opensm	5.20.0.MLNX20240801.ef1f438a-0.1.2407052	InfiniBand Subnet Manager and Subnet Administrator based on OpenSM	GPLv2 BSD
openvswitch	2.17.8-1.2407052	OVS (virtual switch), DPDK based	ASL 2.1 LGPLv3 and SI:
perftest	24.07.0-0.44.g57725f2.2407052	Test suite for performance	BSD 3- Clause

Component	Version	Description	Licens
			v2, or I
rdma-core	2407mlnx52-1.2407052	Implementation of the RDMA verbs	GPLv2 BSD
rivermax	1.51.4	NVIDIA® Rivermax® is an optimized networking SDK for media and data streaming applications	Proprie
rshim	2.0.38-0.gc0f82f3	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	GPLv2
sharp	3.8.0.MLNX20240801.618ff287-1.2407052	Improves the performance of MPI and Machine Learning collective operation by offloading from CPUs and GPUs to the network and eliminating the need to send data multiple times between endpoints	Proprie
sockperf	3.10-0.git5ebd327da983.2407052	Network benchmarking utility over socket API UDP/TCP designed for testing network performance	BSD

Component	Version	Description	Licens
		(latency and throughput)	
spdk	23.01.5-21	SPDK provides a set of tools and libraries for writing high performance, scalable, user-mode storage applications	Proprietary
srp	24.07-OFED.24.07.0.5.2.1	Storage-related driver for SCSI RDMA Protocol initiator	GPLv2
ucx	1.17.0-1.2407052	High-level application-oriented API for high-performance communication over RDMA networks	BSD
virtio-net-controller	24.07.11-1	Virtio-net-controller is a systemd service running on BlueField, with a user interface front-end to manage the emulated virtio-net devices	Proprietary
vma	9.8.60-1	Accelerates latency-sensitive and throughput-demanding TCP and UDP socket-based applications by offloading traffic from the user-space directly to the network interface card	GPLv2 BSD

Component	Version	Description	Licens
		(NIC) or Host Channel Adapter (HCA)	
xlio	3.31.2-1	Boosts the performance of TCP/IP applications based on NGINX (CDN, DoH, etc.) and storage solutions as part of the SPDK	GPLv2 BSD
xpmem	2.7.3-1.2407052	Kernel module to enable inter-process mapping for memory copy in user space	GPLv2 LGPLv2
xpmem-lib	2.7-0.2310055	High-performance inter-process memory sharing	LGPLv2

4. DOCA Packages

Device	Component	Version	Description
Host	DOCA Devel	2.8.0	Software development kit package and tools for developing host software
	DOCA Runtime	2.8.0	Runtime libraries and tools required to run DOCA-based software applications on host
	DOCA Extra	2.8.0	Contains helper scripts (doca-info, doca-kernel-support)
	DOCA OFED	2.8.0	Software stack which operates across all NVIDIA network adapter solutions
	Arm emulated (QEMU) development container	4.8.0	Linux-based BlueField Arm emulated container for developers

Device	Component	Version	Description
Target BlueField DPU (Arm)	BlueField BSP	4.8.0	BlueField image and firmware
	DOCA SDK	2.8.0	Software development kit packages and tools for developing Arm software
	DOCA Runtime	2.8.0	Runtime libraries and tools required to run DOCA-based software applications on Arm

5. Supported Host OS and Features per DOCA-Host Installation Profile

The default operating system included with the **BlueField Bundle** (for **DPU and SuperNIC**) is Ubuntu 22.04.

The supported operating systems on the **host machine** per **DOCA-Host** installation profile are the following:

Note

Only the following generic kernel versions are supported for DOCA local repo package for host installation.

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	Supported DOCA	
			doca-all	doca-network
Alinux 3.2	x86	5.10.134-13.al8.x86_64	✓	✓
Alma 8.5	x86	4.18.0-348.12.2.EL8_5.X86_64	✗	✗
Anolis OS 8.4	aarch64	4.18.0-348.2.1.AN8_4.aarch64	✗	✗

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	Supported DOCA	
			DOCA 5.0	DOCA 6.0
	x86	4.18.0-305.AN8.X86_64	✗	✗
Anolis OS 8.6	aarch64	5.10.134+	✗	✗
	x86	5.10.134+	✗	✗
BCLinux 21.10SP2	aarch64	4.19.90-2107.6.0.0098.oe1.bclinux.aarch64	✗	✗
	x86	4.19.90-2107.6.0.0100.oe1.bclinux.x86_64	✗	✗
BCLinux 22.10	aarch64	5.10.0-153.24.0.100.6.oe2203sp2.bclinux.aarch64	✗	✗
	x86	5.10.0-153.24.0.100.6.oe2203sp2.bclinux.x86_64	✗	✗
CentOS Stream 8	aarch64	4.18.0-552.EL8.AARCH64	✗	✗
	x86	4.18.0-552.el8.x86_64	✗	✗
CentOS Stream 9	aarch64	5.14.0-480.EL9.AARCH64	✗	✗
	x86	5.14.0-480.el9.x86_64	✗	✗
CTyunOS 2.0	aarch64	4.19.90-2102.2.0.0062.ctl2.aarch64	✗	✗
	x86	4.19.90-2102.2.0.0062.ctl2.x86_64	✗	✗
CTyunOS 23.01	aarch64	5.10.0-136.12.0.86.ctl3.aarch64	✓	✓
	x86	5.10.0-136.12.0.86.ctl3.x86_64	✓	✓
Debian 10.8	aarch64	4.19.0-14-arm64	✗	✗
	x86	4.19.0-14-amd64	✓	✓
Debian 10.9	x86	4.19.0-14-amd64	✗	✗
	x86	4.19.0-16-amd64	✗	✗
Debian 10.13	aarch64	4.19.0-21-arm64	✗	✗
	x86	4.19.0-21-amd64	✓	✓
Debian 11.3	aarch64	5.10.0-13-arm64	✗	✗
	x86	5.10.0-13-amd64	✗	✗

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	Supported DOCA	
			1.10	1.11
Debian 12.1	aarch64	6.1.0-10-arm64	✗	✗
	x86	6.1.0-10-amd64	✓	✓
Debian 12.5	aarch64	6.1.0-18-arm64	✓	✗
	x86	6.1.0-18-amd64	✓	✗
EulerOS 2.0 SP9	aarch64	4.19.90- vhulk2006.2.0.h171.eulerosv2r9.aarch64	✗	✗
	x86	4.18.0-147.5.1.0.h269.eulerosv2r9.x86_64	✗	✗
EulerOS 2.0 SP10	aarch64	4.19.90- vhulk2110.1.0.h860.eulerosv2r10.aarch64	✗	✗
	x86	4.18.0- 147.5.2.4.h694.eulerosv2r10.x86_64	✗	✗
EulerOS 2.0 SP11	aarch64	5.10.0- 60.18.0.50.h323.eulerosv2r11.aarch64	✗	✗
	x86	5.10.0- 60.18.0.50.h323.eulerosv2r11.x86_64	✗	✗
EulerOS 2.0 SP12	aarch64	5.10.0- 136.12.0.86.h1032.eulerosv2r12.aarch64	✗	✗
	x86	5.10.0- 136.12.0.86.h1032.eulerosv2r12.x86_64	✗	✗
Kylin 1.0 SP2	aarch64	4.19.90-24.4.v2101.ky10.aarch64	✗	✗
	x86	4.19.90-24.4.v2101.ky10.x86_64	✗	✗
Kylin 1.0 SP3	aarch64	4.19.90-52.22.v2207.ky10.aarch64	✗	✗
	x86	4.19.90-52.22.v2207.ky10.x86_64	✗	✗
Linux Kernel 6.10	aarch64	6.10	✗	✗
	x86		✗	✗
Mariner 2.0	x86	5.15.148.2-2.cm2	✗	✗
Oracle Linux 7.9	x86	5.4.17-2011.6.2.el7uek.x86_64	✗	✗

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	Supported DOCS	
			1	2
Oracle Linux 8.4	x86	5.4.17-2102.201.3.el8uek.x86_64	✗	✗
Oracle Linux 8.6	x86	5.4.17-2136.307.3.1.el8uek.x86_64	✗	✗
Oracle Linux 8.7	x86	5.15.0-3.60.5.1.el8uek.x86_64	✓	✓
Oracle Linux 8.8	x86	5.15.0-101.103.2.1.el8uek.x86_64	✗	✗
Oracle Linux 9.0	x86	5.15.0-0.30.19.el9uek.x86_64	✗	✗
Oracle Linux 9.1	x86	5.15.0-3.60.5.1.el9uek.x86_64	✗	✗
Oracle Linux 9.2	x86	5.15.0-101.103.2.1.el9uek.x86_64	✗	✗
OpenSUSE 15.3	aarch64	-	✗	✗
	x86	5.3.18-150300.59.43-DEFAULT	✗	✗
openEuler 20.03 SP1	aarch64	4.19.90-2012.4.0.0053.OE1.aarch64	✗	✗
	x86	4.19.90-2110.8.0.0119.OE1.X86_64	✗	✗
openEuler 20.03 SP3	aarch64	4.19.90-2112.8.0.0131.oe1.aarch64	✗	✗
	x86	4.19.90-2112.8.0.0131.oe1.x86_64	✗	✗
openEuler 22.03	aarch64	5.10.0-60.18.0.50.oe2203.aarch64	✗	✗
	x86	5.10.0-60.18.0.50.oe2203.x86_64	✗	✗
openEuler 22.03 SP1	x86	5.10.0-136.12.0.86.oe2203sp1.x86_64	✓	✗
Photon OS 3.0	x86	4.19.225-3.ph3	✗	✗
RHEL/CentOS 8.0	aarch64	4.18.0-80.el8.aarch64	✗	✗
	x86	4.18.0-80.el8.x86_64	✗	✗

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	Supported DOCA	
			4.18.0-147.el8	4.18.0-193.el8
RHEL/CentOS 8.1	aarch64	4.18.0-147.el8.aarch64	✗	✗
	x86	4.18.0-147.el8.x86_64	✗	✗
RHEL/CentOS 8.2	aarch64	4.18.0-193.el8.aarch64	✗	✗
	x86	4.18.0-193.el8.x86_64	✓	✓
RHEL/CentOS 8.3	aarch64	4.18.0-240.el8.aarch64	✗	✗
	x86	4.18.0-240.el8.x86_64	✗	✗
RHEL/CentOS 8.4	aarch64	4.18.0-305.el8.aarch64	✗	✗
	x86	4.18.0-305.el8.x86_64	✗	✗
RHEL/CentOS 8.5	aarch64	4.18.0-348.el8.aarch64	✗	✗
	x86	4.18.0-348.el8.x86_64	✗	✗
RHEL/Rocky 8.6	aarch64	aarch644.18.0-372.41.1.el8_6.aarch64	✗	✗
	x86	4.18.0-372.41.1.el8_6.x86_64	✓	✓
RHEL/Rocky 8.7	aarch64	4.18.0-425.14.1.el8_7.aarch64	✗	✗
	x86	4.18.0-425.14.1.el8_7.x86_64	✗	✗
RHEL/Rocky 8.8	aarch64	4.18.0-477.10.1.el8_8.aarch64	✓	✓
	x86	4.18.0-477.10.1.el8_8.x86_64	✓	✓
RHEL/Rocky 8.9	aarch64	4.18.0-513.5.1.el8_9.aarch64	✓	✓
	x86	4.18.0-513.5.1.el8_9.x86_64	✓	✓
RHEL/Rocky 8.10	aarch64	4.18.0-553.el8_10.aarch64	✓	✓
	x86	4.18.0-553.el8_10.x86_64	✓	✓
RHEL/Rocky 9.0	aarch64	5.14.0-70.46.1.el9_0.aarch64	✗	✗
	x86	5.14.0-70.46.1.el9_0.x86_64	✗	✗
RHEL/Rocky 9.1	aarch64	5.14.0-162.19.1.el9_1.aarch64	✗	✗
	x86	5.14.0-162.19.1.el9_1.x86_64	✓	✓

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	Supported DOCA	
			1	2
RHEL/Rocky 9.2	aarch64	5.14.0-284.11.1.el9_2.aarch64	✗	✗
	x86	5.14.0-284.11.1.el9_2.x86_64	✗	✗
RHEL/Rocky 9.3	aarch64	5.14.0-362.8.1.el9_3.aarch64	✗	✗
	x86	5.14.0-362.8.1.el9_3.x86_64	✗	✗
RHEL/Rocky 9.4	aarch64	5.14.0-427.13.1.el9_4.aarch64	✗	✓
	x86	5.14.0-427.13.1.el9_4.x86_64	✗	✓
SLES 15 SP2	aarch64	5.3.18-22-default	✗	✗
	x86	5.3.18-22-default	✗	✗
SLES 15 SP3	aarch64	5.3.18-57-default	✗	✗
	x86	5.3.18-57-default	✗	✗
SLES 15 SP4	aarch64	5.14.21-150400.22-default	✗	✗
	x86	5.14.21-150400.22-default	✗	✗
SLES 15 SP5	aarch64	5.14.21-150500.53-default	✗	✗
	x86	5.14.21-150500.53-default	✗	✗
SLES 15 SP6	x86	6.4.0-150600.21-default	✗	✗
TencentOS 3.3	aarch64	5.4.119-19.0009.39	✗	✗
	x86	5.4.119-19.0009.39	✗	✗
Ubuntu 20.04	aarch64	5.4.0-26-generic	✗	✗
	x86	5.4.0-26-generic	✓	✓
Ubuntu 22.04	aarch64	5.15.0-25-generic	✓	✓
	x86	5.15.0-25-generic	✓	✓
Ubuntu 24.04	aarch64	6.6.0-14-generic	✓	✓
	x86	6.6.0-14-generic	✓	✓
UOS 20.1060	aarch64	5.10.0-46.uel20.aarch64	✗	✗
	x86	5.10.0-46.uel20.x86_64	✗	✗

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	Supported DOCA-OFED	
			DOCA-OFED	DOCA-OFED
UOS 20.1060a	aarch64	5.10.0-46.uelc20.aarch64	✗	✗
	x86	5.10.0-46.uelc20.x86_64	✗	✗

6. DOCA-OFED Version Interoperability

This section reflects which versions were tested and verified for multi-version environments (i.e., environments with more than one doca-oped version on host servers).

Target Version	Versions Verified for Interoperability
24.07-1.x.x.x July 2024	24.04-0.7.0.0 - DOCA-OFED Profile
	5.8-5.1.1.2 LTS

7. BF-Bundle (BFB) Version Upgrade/Downgrade

The following table provides a matrix for the supported upgrade/downgrade of BFBs across different versions.

Version	Upgrade to	Downgrade to
1.5.0	2.0.2; 2.2.0; 1.5.1; 1.5.2; 1.5.3	1.4.0; 1.3.0
1.5.1	1.5.2	1.5.0
1.5.2	1.5.3	1.5.1; 1.5.0
1.5.3	N/A	1.5.2; 1.5.0
2.0.2	2.2.0; 2.5.0	1.5.0; 1.4.0
2.2.0	2.5.0; 2.6.0	N/A
2.2.1	2.5.0; 2.6.0	N/A
2.5.0	2.5.1; 2.6.0	2.2.1 for BlueField-3; 2.2.0 for BlueField-2
2.5.1	2.5.2	2.5.0
2.5.2	N/A	2.5.1; 2.5.0

Version	Upgrade to	Downgrade to
2.6.0	2.7.0	2.5.0; 2.2.1 for BlueField-3; 2.2.0 for BlueField-2
2.7.0	2.8.0	2.6.0; 2.5.0; 2.2.1 for BlueField-3; 2.2.0 for BlueField-2
2.8.0	N/A	2.7.0; 2.6.0; 2.5.0

8. Supported DOCA Version Upgrade Using Standard Linux Tools on BlueField

Version	Upgrade to
2.5.0	2.5.1; 2.6.0; 2.7.0; 2.8.0
2.5.1	2.5.2
2.5.2	N/A
2.6.0	2.7.0; 2.8.0
2.7.0	2.8.0

9. API Changes

Note

The [old DOCA Comm Channel API](#) will be deprecated in DOCA 2.9.0.

Library	Change Description
doca_comch	<ul style="list-style-type: none"> Changed features <ul style="list-style-type: none"> API function name changes API function parameter and return value changes
doca_dma	<ul style="list-style-type: none"> Added features <ul style="list-style-type: none"> Enable exporting DMA to GPU
doca_dpa	<ul style="list-style-type: none"> Added features

Library	Change Description
	<ul style="list-style-type: none"> ◦ Add multi-GVMI support (i.e., run DOCA DPA RDMA on VF while DOCA DPA created on PF)
doca_common	<ul style="list-style-type: none"> • Added features <ul style="list-style-type: none"> ◦ Bitfield support ◦ Expandable <code>doca_buf_inventory</code> ◦ Batching support (group tasks and flash explicitly to hardware) ◦ Set <code>doca_pe</code> (progress-engine) affinity • Changed features <ul style="list-style-type: none"> ◦ Imported <code>doca_mmap</code> (to DPU) can be exported to (remote) RDMA
doca_compress	<ul style="list-style-type: none"> • Removed features <ul style="list-style-type: none"> ◦ Decompress LZ4
doca_eth	<ul style="list-style-type: none"> • Added features <ul style="list-style-type: none"> ◦ Ability to extend (i.e., increase) number of allocated tasks ◦ Control notification moderation (once in <code>n</code> events or <code>time</code>) • Changed features <ul style="list-style-type: none"> ◦ Parameter order in: <code>doca_eth_rxq_task_recv_allocate_init</code>
doca_gpunetio	<ul style="list-style-type: none"> • Added features <ul style="list-style-type: none"> ◦ Support <code>doca_buf</code> on GPU (<code>doca_gpu_buf</code>) ◦ Support dma <code>operations</code> GPU ↔ DPU/host • Changed Features <ul style="list-style-type: none"> ◦ RDMA API changes
doca_pcc	<ul style="list-style-type: none"> • Added features <ul style="list-style-type: none"> ◦ More debug/dump APIs ◦ Performance enhancements (e.g., inline functions) • Changed features <ul style="list-style-type: none"> ◦ Structure of <code>cc_event</code> – Added support for future hardware (placeholder)

10. Device Definition

The supported adapter cards are specified as follows:

Supported Cards	Description
All HCAs	Supported in the following adapter cards unless specifically stated otherwise: ConnectX-4/ConnectX-4 Lx/ConnectX-5/ConnectX-6/ConnectX-6 Dx/ConnectX-6 Lx/ConnectX-7/BlueField-2/BlueField-3
ConnectX-6 Dx and above	Supported in the following adapter cards unless specifically stated otherwise: ConnectX-6 Dx/ConnectX-6 Lx/ConnectX-7/BlueField-2/BlueField-3
ConnectX-6 and above	Supported in the following adapter cards unless specifically stated otherwise: ConnectX-6/ConnectX-6 Dx/ConnectX-6 Lx/ConnectX-7/BlueField-2/BlueField-3
ConnectX-5 and above	Supported in the following adapter cards unless specifically stated otherwise: ConnectX-5/ConnectX-6/ConnectX-6 Dx/ConnectX-6 Lx/ConnectX-7/BlueField-2/BlueField-3
ConnectX-4 and above	Supported in the following adapter cards unless specifically stated otherwise: ConnectX-4/ConnectX-4 Lx/ConnectX-5/ConnectX-6/ConnectX-6 Dx/ConnectX-6 Lx/ConnectX-7/BlueField-2/BlueField-3

11. Unsupported Functionalities/Features/NICs

The following are the unsupported functionalities/features/NICs in the current version:

- RDMA experimental verbs library (mlnx_lib)
- CIFS (Common Internet File System) module installation
- Relational Database Service (RDS)
- mthca InfiniBand driver
- Ethernet IPoIB (eIPoIB)
- InfiniBand Connected transport service
- IPSec over bond for crypto offload

Changes and New Features

1. New Features and Updates

Note

DOCA 2.8.0 is a mandatory update release for all customers and projects with BlueField-3 DPU or SuperNIC when used in NIC mode with Arm cores disabled. This version fixes an eMMC clock toggling loop issue after boot is completed.

Note

The October '24 LTS release will be the final software version to support ConnectX-4 device. Starting January '25, Connect-X 4 will no longer be supported by future DOCA-Host releases.

Note

BlueField-3 networking platforms are required to use DOCA-Host as the host driver. MLNX_OFED does not support BlueField-3 devices.

- Spectrum-X 1.1 – SuperNIC Enhancements and Host Telemetry – OTLP streaming protocol
- DOCA-Flow and OVS-DOCA enhancements – Hitless upgrade/restart, micro-segmentation, "send-to-kernel" switch mode, "Basic pipe" resize, sFlow support

(monitoring, debugging)

- Added alpha support for SNAP Virtio-FS file system emulation to the early access NGC service container
- Added beta support for DOCA Management Service (DMS) – systemd service in DOCA for Host package
- DOCA DPA resource allocation optimization – Allocating DPA compute resources to multiple apps
- DOCA Core – Added support for L3 cache invalidation and task batching submission/completion
- DOCA reference applications code and DOCA libs sample code is now provided under BSD-3 open-source license
- New DOCA Comch library (Comm Channel) GA, will replace the previous DOCA Comm Channel library which is scheduled to be deprecated in the next release (Oct '24). See [DOCA Comch](#) for details.
- BFB update – Added support for setting BMC password
- Added new BF-Bundle package format, `.iso`, in addition to `.bfb`

Bug Fixes in This Version

1. DOCA Bug Fixes

Ref #	Issue
3928479	Description: Users may encounter an error in "dmesg" when unplugging an emulated PCIe device.
	Keyword: DevEmu
	Reported in version: 2.7.0
3881941	Description: When working with RShim 2.0.28, PCIe host crash may rarely occur at the beginning of BFB push after the Arm reset.
	Keyword: RShim; driver
	Reported in version: 2.7.0
3882794	Description: When working with <code>doca_pcc_np</code> context, the return value from the API <code>doca_pcc_get_max_num_threads()</code> is incorrect. The function has an output parameter that indicates the maximum number of threads allowed for a <code>doca_pcc_np</code> context. The correct value that the library expects is 16 instead of the returned 64.
	Keyword: PCC; threads
	Reported in version: 2.7.0
3840230	Description: Order of cores specified in <code>--core-list</code> is not respected. Cores are picked in ascending order instead.
	Keyword: DOCA Bench
	Reported in version: 2.7.0
3849701	Description: DOCA Comch tests cannot be launched from BlueField side .
	Keyword: DOCA Bench; DOCA Comch
	Reported in version: 2.7.0
3857097	Description: DOCA RDMA tests cannot be launched from BlueField side.
	Keyword: DOCA Bench; DOCA RDMA

Ref #	Issue
	Reported in version: 2.7,0
3857095	Description: Send tasks on DOCA RDMA may fail.
	Keyword: DOCA Bench; DOCA RDMA; send
	Reported in version: 2.7,0
3859823	Description: Multi-threaded tests using DOCA Comch may hang or emit an infinite amount of log messages. Single-threaded tests are less likely to cause this issue.
	Keyword: DOCA Bench; DOCA Comch
	Reported in version: 2.7.0
3872654	Description: And i ssue occurs when submitting tasks with DOCA SHA with the following error.
	<pre>[DOCA][ERR][doca_pe.cpp:177][task_submit] Task 0xaaaaf4865bf0: Failed to submit task: task is already submitted</pre>
	Keyword: DOCA Bench
	Reported in version: 2.7,0
3869639	Description: Users c annot use <code>--job-output-buffer-size 0</code> when using remote output memory (<code>--use-remote-output-buffers</code>).
	Keywords: DOCA Bench
	Reported in version: 2.7,0
3886315	Description: To reset or shut down the BlueField Arm, it is mandatory to specify the <code>--sync 0</code> argument with reset level 1 and reset type 3 or 4. For example:
	<pre>m1xfwreset -d <device> -l 1 -t 4 --sync 0 r</pre>
	Keyword: Arm; shutdown
	Reported in version: 2.7.0

Ref #	Issue
3957990	Description: Sending a malformed UDP packet with VXLAN configuration causes OVS-DOCA to crash.
	Keyword: OVS-DOCA; encap; crash
	Reported in version: 2.7.0
3994490	Description: Malformed packets cause OVS to crash when performing encapsulation.
	Keyword: Openvswitch
	Reported in version: 2.7.0
3949342	Description: NVQual fails due to low line rate.
	Keyword: NVQual
	Reported in version: 2.7.0
3960883	Description: If working with 2 different NICs with the same app, encryption can occur on the wrong port.
	Keyword: PSP gateway
	Reported in version: 2.6.0
3546202	Description: After rebooting a BlueField-3 DPU running Rocky Linux 8.6 BFB, the kernel log shows the following error:
	<pre>[3.787135] mlxbf_gige MLNXBF17:00: Error getting PHY irq. Use polling instead</pre>
	This message indicates that the Ethernet driver will function normally in all aspects, except that PHY polling is enabled.
	Keywords: Linux; PHY; kernel
	Reported in version: 2.2.0

2. BSP Bug Fixes

Ref #	Issue Description
3836598	Description: BIOS attribute in Redfish schema "SPCR UART" does not comply with expected format.
	Keyword: Redfish; BIOS

Ref #	Issue Description
	Reported in version: 4.6.0
4035187	Description: When eMMC is stuck, Linux tries to recover it including resetting the eMMC itself and re-initializing it.
	Keyword: eMMC stuck
	Reported in version: 4.8.0

3. BMC Bug Fixes

Ref #	Issue
4047689	Description: While running the reprovisioning script from BMC, the RShim boot device appears to be busy which causes the script to fail without completing the process.
	Discovered in version: 24.07
4146640	Description: In the event of a server reboot, the BMC may boot before the host and take control of the RShim before the host.
	Discovered in version: 24.07
3991930	Description: The reported dump entry creation date is not initialized properly and reports the default system date <code>1970-01-01T00:28:43.991149+00:00</code> when creating a dump entry using the LogService on the Redfish interface.
	Discovered in version: 24.07
4064371	Description: The BMC dump collection is missing the <code>varfilelist.log</code> and <code>slabinfo.log</code> files.
	Discovered in version: 24.07

Known Issues

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

Reference	Description
4032924	Description: When upgrading to DOCA 2.8.0 on RPM-based OSES, a conflict between <code>strongswan-bf</code> or <code>libreswan</code> and <code>strongSwan</code> may occur.
	Workaround: Before upgrading, delete <code>strongswan-bf</code> and <code>libreswan</code> : <pre>yum remove strongswan-bf strongswan-swanctl libreswan</pre>
	Keyword: <code>strongSwan</code> ; upgrade
	Reported in version: 2.8.0
4035553	Description: <code>oper_sample_period</code> does not always reflect the correct sample period. In some cases, it will reflect the <code>admin_sample_period</code> instead.
	Workaround: N/A
	Keyword: Core
	Reported in version: 2.8.0
4023257	Description: If RDMA samples are compiled with memory sanitizer enabled, "read memory leak" errors are printed when running the samples with the RDMA CM flag and when running the client before the server.
	Workaround: Make sure to start the RDMA Server before RDMA Client.
	Keyword: DOCA RDMA; samples
	Reported in version: 2.8.0
4021752 4021748	Description: In all RDMA samples, if an error occurs in any of the following functions: <ul style="list-style-type: none"> Exporting RDMA/MMAP/Sync event Connecting RDMA Writing or reading the descriptors

Reference	Description
	An error is printed but the sample resumes and might: <ul style="list-style-type: none"> 1. Fail later, or be in busy-wait state indefinitely; and/or 2. Result in access to an unknown address, causing an address sanitizer violation.
	Workaround for 1: Either: <ul style="list-style-type: none"> • Follow the error logs to verify no errors occurred in the relevant function. And if it did, stop the sample. • Fix the issue locally.
	Workaround for 2: The mentioned address sanitizer violation shall be ignored in case of an error in a relevant function.
	Keyword: DOCA RDMA; samples
	Reported in version: 2.8.0
3961940	Description: OVS-DOCA connection tracking with E2E enabled is not supported.
	Workaround: N/A
	Keyword: OVS-DPDK; connection tracking; E2E
	Reported in version: 2.8.0
3989851	Description: A DOCA Flow pipe has multiple actions. When the action <code>idx</code> is not 0 and it has a shared endecap action, a crash occurs when attempting to create an entry.
	Workaround: N/A
	Keyword: DOCA Flow
	Reported in version: 2.8.0
3988904	Description: Failure to create a control entry with shared endecap action.
	Workaround: N/A
	Keyword: DOCA Flow
	Reported in version: 2.8.0
3886674	Description: Installing <code>doca-all</code> and other DOCA metapackages does not install the <code>m1nx-nvme</code> driver.

Reference	Description
	<p>Workaround: <code>mlnx-nvme</code> is only needed for NVMe-over-RDMA remote storage support. If you wish to install it, add the <code>mlnx-nvme</code> package to the install command.</p> <ul style="list-style-type: none"> On RHEL: <pre>apt install doca-all mlnx-nvme-modules</pre> On Ubuntu: <pre>dnf install doca-all-kmod-mlnx-nvme</pre> <p>Keyword: NVMe; DOCA profile</p> <p>Reported in version: 2.7.0</p>
3885930	<p>Description: When installing DOCA-Host on a system using NVMe storage (typically local NVMe disk), and the script <code>doca-kernel-support</code> is used to rebuild and install kernel modules, unloading the <code>mlx5</code> drivers is only possible after also unmounting the NVMe storage, which would typically necessitate a reboot.</p> <p>Workaround: N/A</p> <p>Keyword: NVMe; doca-kernel-support; DOCA for host</p> <p>Reported in version: 2.7.0</p>
3837255	<p>Description: When running Arm shutdown from the host OS it is expected to get the message <code>-E- Failed to send Register MRSI</code>. This message should be ignored.</p> <p>Workaround: Wait 2 more minutes before rebooting the host. Before proceeding with host OS reboot, it is recommended to query the operational state of the BlueField Arm cores from the BlueField BMC to verify that shutdown state has been reached. Run the following command:</p> <pre>ipmitool -C 17 -I lanplus -H <bmc_ip> -U root -P <password> raw 0x32 0xA3</pre>

Reference	Description
	Expected output is "06".
	Keyword: Host OS; reboot; error
	Reported in version: 2.7.0
3844705	Description: In OpenEuler 20.03, the Linux Kernel version 4.19.90 is affected by an issue that impacts the discard/trim functionality for the BlueField eMMC device which may cause degraded performance of the BlueField eMMC over time.
	Workaround: Upgrade to Linux Kernel version 5.10 or later.
	Keyword: eMMC discard; trim functionality
	Reported in version: 2.7.0
3877725	Description: During BFB installation in NIC mode on BlueField-3, too much information is added into RShim log which fills it, causing the Linux installation progress log to not appear in the RShim log.
	<pre>echo "DISPLAY_LEVEL 2" > /dev/rshim0/misc cat /dev/rshim0/misc</pre>
	Workaround: Monitor the BlueField-3 Arm's UART console to check whether BFB installation has completed or not for NIC mode.
	<pre>[13:58:39] INFO: Installation finished ... [14:01:53] INFO: Rebooting...</pre>
	Keyword: NIC mode; BFB install
Reported in version: 2.7.0	
3855702	Description: Trying to jump from a steering level in the hardware to a lower level using software steering is not supported on rdma-core lower than 48.x.
	Workaround: N/A
	Keyword: RDMA; SWS
	Reported in version: 2.7.0

Reference	Description
3855485	<p>Description: When enabling the <code>PCI_SWITCH_EMULATION_ENABLE</code> NVconfig, the mlx devices, and potentially the RShim devices disappear. Also, looking at the kernel logs using <code>dmesg</code> shows the following messages:</p> <pre data-bbox="332 359 1461 720"> pci 0000:29:00.0: BAR 0: no space for [mem size 0x0200 0000 64bit pref] pci 0000:29:00.0: BAR 2: no space for [mem size 0x0080 0000 64bit pref] ... </pre> <p>Workaround: N/A</p> <p>Keyword: NVconfig; RShim; dmsg</p> <p>Reported in version: 2.7.0</p>
3831230	<p>Description: In OpenEuler 20.03, the Linux Kernel version 4.19.90 is affected by an issue that impacts the discard/trim functionality for BlueField eMMC device which may cause degraded performance of BlueField eMMC over time.</p> <p>Workaround: Upgrade to Linux Kernel version 5.10 or later.</p> <p>Keyword: eMMC discard; trim functionality</p> <p>Reported in version: 2.7.0</p>
3743879	<p>Description: <code>mlxfwreset</code> could timeout on servers where the RShim driver is running and INTx is not supported. The following error message is printed:</p> <pre data-bbox="332 1381 1461 1522"> BF reset flow encountered a failure due to a reset state error of negotiation timeout . </pre> <p>Workaround: Set <code>PCIE_HAS_VFIO=0</code> and <code>PCIE_HAS_UIO=0</code> in <code>/etc/rshim.conf</code> and restart the RShim driver. Then re-run the <code>mlxfwreset</code> command.</p> <p>If host Linux kernel lockdown is enabled, then manually unbind the RShim driver before <code>mlxfwreset</code> and bind it back after <code>mlxfwreset</code>:</p> <pre data-bbox="332 1774 1461 1944"> echo "DROP_MODE 1" > /dev/rshim0/misc mlxfwreset <arguments> </pre>

Reference	Description
	<pre>echo "DROP_MODE 0" > /dev/rshim0/misc</pre>
	Keyword: Timeout; mlxfwreset; INTx
	Reported in version: 2.7.0
3665070	Description: Virtio-net controller fails to load if <code>DPA_AUTHENTICATION</code> is enabled.
	Workaround: N/A
	Keyword: Virtio-net; DPA
	Reported in version: 2.5.0
3678069	Description: If using BlueField with NVMe and mmcbld and configured to boot from mmcblk, users must create <code>bf.cfg</code> file with <code>device=/dev/mmcblk0</code> , then install the <code>*.bfb</code> as normal.
	Workaround: N/A
	Keyword: NVMe
	Reported in version: 2.5.0
3680538	Description: When using strongSwan or OVS-IPsec as explained in the NVIDIA BlueField DPU BSP , the IPsec Rx data path is not offloaded to hardware and occurs in software running on the Arm cores. As a result, bandwidth performance is substantially low.
	Workaround: N/A
	Keyword: IPsec
	Reported in version: 2.5.0
N/A	Description: Execution unit partitions are still not implemented and would be added in a future release.
	Workaround: N/A
	Keyword: EU tool
	Reported in version: 2.5.0
3666160	Description: Installing BFB using <code>bfb-install</code> when <code>mlxconfig PF_TOTAL_SF > 1700</code> , triggers server reboot immediately.

Reference	Description
	<p>Workaround: Change <code>PF_TOTAL_SF</code> to 0, perform a graceful shutdown, power cycle, then installing BFB.</p> <p>Keyword: SF; <code>PF_TOTAL_SF</code>; BFB installation</p> <p>Reported in version: 2.2.1</p>
3594836	<p>Description: When enabling Flex IO SDK tracer at high rates, a slow-down in processing may occur and/or some traces may be lost.</p> <p>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.</p> <p>Keyword: Tracer FlexIO</p> <p>Reported in version: 2.2.1</p>
3592080	<p>Description: When using UEK8 on the host in DPU mode, creating a VF on the host consumes about 100MB memory on BlueField</p> <p>Workaround: N/A</p> <p>Keyword: UEK; VF</p> <p>Reported in version: 2.2.1</p>
3546202	<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>
3566042	<p>Description: Virtio hotplug is not supported in GPU-HOST mode on the NVIDIA Converged Accelerator.</p> <p>Workaround: N/A</p> <p>Keyword: Virtio; Converged Accelerator</p>

Reference	Description
	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 BlueField, 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
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
	Keywords: mlxfwreset
	Reported in version: 2.2.0
3306489	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".
	Workaround: Add <code>intel_idle.max_cstate=1</code> entry to the kernel command line.
	Keywords: Longevity; mlxfwreset; DPU reboot
	Reported in version: 2.2.0
3538486	Description: When removing LAG configuration from BlueField, a kernel warning for <code>uverbs_destroy_ufile_hw</code> is observed if virtio-net-controller is still running.
	Workaround: Stop virtio-net-controller service before cleaning up bond configuration.
	Keywords: Virtio-net; LAG
	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:

Reference	Description
	<p>1. Run:</p> <pre data-bbox="412 254 1463 457">echo 0 > /sys/bus/platform/drivers/mlxbf-bootctl/large_icm</pre> <p>2. Reboot Arm.</p> <p>Keyword: BlueField-3; downgrade</p> <p>Reported in version: 2.2.0</p>
3462630	<p>When trying to perform a PXE installation when UEFI Secure Boot is enabled, the following error messages may be observed:</p> <pre data-bbox="331 758 1463 961">error: shim_lock protocol not found. error: you need to load the kernel first.</pre> <p>Workaround: Download a Grub EFI binary from the Ubuntu website. For further information on Ubuntu UEFI Secure Boot PXE Boot, please visit Ubuntu's official website.</p> <p>Keyword: PXE; UEFI Secure Boot</p> <p>Reported in version: 2.0.2</p>
3448841	<p>Description: While running CentOS 8.2, switchdev Ethernet BlueField runs in "shared" RDMA net namespace mode instead of "exclusive".</p> <p>Workaround: Use <code>ib_core</code> module parameter <code>netns_mode=0</code>. For example:</p> <pre data-bbox="331 1440 1463 1644">echo "options ib_core netns_mode=0" >> /etc/modprobe.d/mlnx-bf.conf</pre> <p>Keyword: RDMA; isolation; Net NS</p> <p>Reported in version: 2.0.2</p>
2706803	<p>Description: When an NVMe controller, SoC management controller, and DMA controller are configured, the maximum number of VFs is limited to 124.</p>

Reference	Description
	Workaround: N/A
	Keyword: VF; limitation
	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
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

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 (“NVIDIA”) makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes 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 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 and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

© Copyright 2025, NVIDIA. PDF Generated on 10/09/2025