



NVIDIA DOCA

Release Notes

Table of Contents

Chapter 1. Introduction.....	1
Chapter 2. Issues Fixed.....	2
Chapter 3. Installation Notes.....	3
3.1. DOCA Packages.....	3
3.2. Supported Operating System Distributions.....	4
Chapter 4. Technical Support.....	5
Chapter 5. Known Issues.....	6

Chapter 1. Introduction

DOCA 1.5.3 is an LTS update to DOCA 1.5.0 which includes bug fixes from DOCA 1.5.2.

Chapter 2. Issues Fixed

This version introduces the following bug fixes.

Reference	Description
3607549	Description: While running <code>iperf</code> on virtio hotplug devices with a guest OS running CentOS kernel 3.10, performing unplug may results in guest kernel stuck and failed to do unplug due to timeout.
	Reported in version: DOCA 1.5.2
3518977	Description: (CVE) ICMP functionality in the Linux Kernel exposes UDP source port randomization (such as DNS servers) to DNS cache poisoning attacks.
	Reported in version: DOCA 1.5.2
3590343	Description: BMC failed login password change loop.
	Reported in version: BMC 2.8.2-46
3586685	Description: App Shield initialization fix on Windows host\VM environment.
	Reported in version: DOCA 2.0

Chapter 3. Installation Notes

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

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

3.1. DOCA Packages

Device	Component	Version	Description
Host	DOCA SDK	1.5.3	Software development kit package for developing host software
	DOCA Runtime	1.5.3	Runtime libraries required to run DOCA-based software applications on host
	DOCA Tools	1.5.3	Tools for developers and administrators on host
	Arm emulated (QEMU) development container	3.9.7	Linux-based BlueField Arm emulated container for developers
Target BlueField-2 DPU (Arm)	BlueField BSP	3.9.7	BlueField image and firmware
	DOCA SDK	1.5.3	Software development kit packages for developing Arm software
	DOCA Runtime	1.5.3	Runtime libraries required to run DOCA-based software applications on Arm
	DOCA Tools	1.5.3	Tools for developers and administrators for Arm target

3.2. 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:

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

Host Operation System	Kernel Support	Arch Support
CentOS 7.6	4.14.0-115.el7a.aarch64	aarch64
	3.10.0-957.el7.x86_64	x86
CentOS 8.0	4.18.0-80.el8.x86_64	
CentOS 8.2	4.18.0-193.el8.x86_64	
RHEL 7.6	3.10.0-957.el7.x86_64	
RHEL 8.0	4.18.0-80.el8.x86_64	
RHEL 8.2	4.18.0-193.el8.x86_64	
Rocky 8.6	4.18.0-372.41.1.el8.x86_64	
Ubuntu 18.04	4.15.0-20-generic	
Ubuntu 20.04	5.4.0-26-generic	
Ubuntu 22.04	5.15.0-52-generic	
Debian 10.8	4.19.0-14-amd64	

Chapter 4. Technical Support

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

- ▶ E-mail: enterprisesupport@nvidia.com
- ▶ 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
3722236	Description: Calling <code>doca_comm_channel_ep_listen()</code> on Rocky or CentOS/RHEL 8.2 BFB fails, causing <code>doca_comm_channel</code> to not work.
	Workaround: N/A
	Keywords: DOCA Comm Channel
	Reported in version: 1.5.3
3374176	Description: Occasionally, an SF representor name fetched from virtio-net does not follow the naming convention due to udev.
	Workaround: Use IP command to rename the SF representor accordingly.
	Keywords: SF representor; virtio-net
	Reported in version: 1.5.2
3282113	Description: Some DPUs experience an issue with the clock settings after installing a BlueField OS in an HBN setting in which the date reverts back to "Thu Sep 8, 2022".
	Workaround: Issue the following commands to enable NTP and sync the clock:
	<pre>sudo apt update sudo apt install ntp sudo systemctl stop ntp sudo systemctl disable ntp sudo systemctl enable ntp@mgmt sudo systemctl start ntp@mgmt sudo systemctl daemon-reload sudo systemctl status ntp@mgmt</pre>
	Keyword: HBN; time; NTP
Reported in version: 1.5.1	
3274589	Description: Comm channel fails to initiated in Rocky Linux.
	Workaround: N/A
	Keyword: Comm channel; Rocky
	Reported in version: 1.5.1
3270534	Description: The Security Gateway reference application fails to forward traffic when used with a ConnectX setup in full-offload mode.
	Workaround: Change the number of queues passed to <code>doca_flow_init()</code> at the application level instead of using the current value.
	Keyword: Security Gateway; IPSec; flow

Reference	Description
	Reported in version: 1.5.1
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> <p>Workaround: N/A</p> <p>Keyword: RegEx; hugepages</p> <p>Reported in version: 1.5.1</p>
3240785	<p>Description: DTS might fail to connect to DPE if started after DTS is already running.</p> <p>Workaround: Start the DTS container only after starting DPE.</p> <p>Keyword: DTS; DPE; BlueMan</p> <p>Reported in version: 1.5.0</p>
3250391	<p>Description: <code>dpgk_queues_and_ports_init</code> fails when given 0 DPDK ports.</p> <p>Workaround: N/A</p> <p>Keyword: <code>doca_dpi_grpc</code> server; DPDK</p> <p>Reported in version: 1.5.0</p>
3240153	<p>Description: DOCA kernel support only works on a non-default kernel.</p> <p>Workaround: N/A</p> <p>Keyword: Kernel</p> <p>Reported in version: 1.5.0</p>
3239630	<p>Description: FlexIO sample <code>flexio_rpc</code> fails to compile.</p> <p>Workaround: Add the following at line 34 in the file <code>/opt/mellanox/doca/samples/flexio/flexio_rpc/meson.build</code>:</p> <pre>source_file = meson.current_source_dir() + '/device/ flexio_rpc_device.c'</pre> <p>Keyword: FlexIO; DOCA samples; <code>flexio_rpc</code></p> <p>Reported in version: 1.5.0</p>
3217627	<p>Description: The <code>doca_devinfo_rep_list_create</code> API returns success on the host instead of "Operation not supported".</p> <p>Workaround: N/A</p> <p>Keyword: DOCA core; InfiniBand</p> <p>Reported in version: 1.5.0</p>
3048250	<p>Description: When configuring the DPU to operate in NIC Mode, the following parameters must be set to default (i.e., =0): <code>HIDE_PORT2_PF</code>, <code>NVME_EMULATION_ENABLE</code>, and <code>VIRTIO_NET_EMULATION_ENABLE</code>.</p> <p>Workaround: N/A</p> <p>Keyword: DPU operation mode</p> <p>Reported in version: 1.3.0</p>

Reference	Description
3017202	Description: Due to disabled backend foundation units, some commands show 500 INTERNAL SERVER ERROR/ 404 NOT FOUND. These commands are related to features or sub-systems which are not supported on HBN.
	Workaround: N/A
	Keyword: HBN; unsupported NVUE commands
	Reported in version: 1.3.0
2821785	Description: MAC addresses are not learned in the hardware but only in software. This may affect performance in pure L2 unicast traffic. This should not affect performance of IPv4/IPv6 traffic or L2 control traffic (i.e., STP, LLDP).
	Workaround: N/A
	Keyword: HBN
	Reported in version: 1.3.0
2828838	Description: NetworkManager and other services not directly related to HBN may display the following message in syslog: "netlink: read: too many netlink events. Need to resynchronize platform cache" The message has no functional impact and may be ignored.
	Workaround: N/A
	Keyword: HBN
	Reported in version: 1.3.0
3168683	Description: If many interfaces are participating in EVPN/routing, it is possible for the routing process to run out of memory.
	Workaround: Have a maximum of 8 VF interfaces participating in routing/VXLAN.
	Keyword: HBN; routing; memory
	Reported in version: 1.2.0
3219539	Description: TC rules are programmed by OVS to map uplink and host representor ports to HBN service. These rules are ageable and can result in packets needing to get software forwarded periodically to refresh the rules.
	Workaround: The timeout value can be adjusted by changing the OVS parameter <code>other_config : max-idle</code> as documented here . The shipped default value is 10000ms (10s).
	Keyword: HBN; SFC; aging
	Reported in version: 1.2.0
3184745	Description: The command <code>nv show interface <intf> acl</code> does not show correct information if there are multiple ACLs bound to the interface.
	Workaround: Use the command <code>nv show interface <intf></code> to view the ACLs bound to an interface.
	Keyword: HBN; ACLs
	Reported in version: 1.2.0
3158934	Description: Deleting an NVUE user by removing their password file and restarting the <code>decrypt-user-add</code> service on the HBN container does not work.
	Workaround: Either respawn the container after deleting the file, or delete the password file corresponding to the user by running <code>userdel -r username</code> .

Reference	Description
	Keyword: HBN; user deletion
	Reported in version: 1.2.0
3191433	Description: ECMP selection for the underlay path uses the ingress port and identifies uplink ports via round robin. This may not result in uniform spread of the traffic.
	Workaround: N/A
	Keyword: HBN; ECMP
	Reported in version: 1.2.0
3185003	Description: When a packet is encapsulated with a VXLAN header, it adds extra bytes which may cause the packet to exceed the MTU of link. Typically, the packet would be fragmented but its silently dropped and no fragmentation happens.
	Workaround: Make sure that the MTU on the uplink port is always 50 bytes more than host ports so that even after adding VXLAN headers, ingress packets do not exceed the MTU.
	Keyword: HBN; MTU; VXLAN
	Reported in version: 1.2.0
3184905	Description: On VXLAN encapsulation, the DF flag is not propagated to the outer header. Such a packet may be truncated when forwarded in the kernel, and it may be dropped when hardware offloaded.
	Workaround: Make sure that the MTU on the uplink port is always 50 bytes more than host ports so that even after adding VXLAN headers, ingress packets do not exceed the MTU.
	Keyword: HBN; VXLAN
	Reported in version: 1.2.0
3188688	Description: When stopping the container using the command <code>crictl stop</code> an error may be reported because the command uses a timeout of 0 which is not enough to stop all the processes in the HBN container.
	Workaround: Pass a timeout value when stopping the HBN container by running: <code>crictl stop --timeout 60 <hbn-container></code>
	Keyword: HBN; timeout
	Reported in version: 1.2.0
3129749	Description: The same ACL rule cannot be applied in both the inbound and outbound direction on a port.
	Workaround: N/A
	Keyword: HBN; ACLs
	Reported in version: 1.2.0
3126560	Description: The system's time zone cannot be modified using NVUE in the HBN container.
	Workaround: The timezone can be manually changed by symlinking the <code>/etc/localtime</code> file to a binary time zone's identifier in the <code>/usr/share/zoneinfo</code> directory. For example: <code>sudo ln -sf /usr/share/zoneinfo/GMT /etc/localtime</code>
	Keyword: HBN; time zone; NVUE

Reference	Description
	Reported in version: 1.2.0
3118204	Description: Auto-BGP functionality (where the ASN does not need to be configured but is dynamically inferred by the system based on the system's role as a leaf or spine device) is not supported on HBN.
	Workaround: If BGP is configured and used on HBN, the BGP ASN must be manually configured.
	Keyword: HBN; BGP
	Reported in version: 1.2.0
3233088	Description: Since checksum calculation is offloaded to the hardware (not done by the kernel), it is expected to see an incorrect checksum in the tcpdump for locally generated, outgoing packets. BGP keepalives and updates are some of the packets that show such incorrect checksum in tcpdump.
	Workaround: N/A
	Keyword: HBN; BGP
	Reported in version: 1.2.0

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

© 2024 NVIDIA Corporation & affiliates. All rights reserved.