



## **Changes and New Features**

# Table of contents

Customer Affecting Changes	7
Declared Unsupported Features	12

# DOCA Framework New Features and Updates

- Support for **ConnectX-8 SuperNIC** - PREL
- Support for **InfiniBand Quantum-X800** (XDR)
- Support for **Spectrum-X 1.3** Reference Architecture
- **DPF** Beta release – Orchestration of DPU and services with K8s
- **DOCA-Flow**
  - Support GTP-U PDU Session Container (PSC) extension header for 5G networks
  - A more robust API/callback to get counters data after a crash
  - Due to infrastructure changes in DOCA Flow, some functionality is impacted. **Please see affecting changes in [Affecting Changes](#) page.**
- **HBN** – Tenant encapsulated packets offload, CoPP enhancements, Stateful NAT+DNAT+PAT support
- **DOCA Ngauge** tool – New tool for debugging network performance with high frequency sampling
- **OVS-DOCA** – Enable creation of optimized anchor pipe
- **DMS** – [mlxconfig reset via DMS](#), [set AR capable packets via DSCP instead of the reserved BTH](#)
- **DOCA P4 Runtime and DOCA P4 Dev Services** Beta release - Write, Compile, Load and Debug your DOCA pipeline with the DOCA P4 language and Tools
- **Congestion control** –
  - Query DPA HW Version and select matching DPA Program, DPACC multiple target support
- **DPU RAS** – Added support for monitoring of BlueField DPU hardware alarms by recording faults of CPU and memory into DPU BMC SEL

# DOCA-Host and DOCA Drivers New Features and Updates

## Note

NVIDIA® BlueField®-3 networking platforms are required to use DOCA-Host as the host driver. MLNX\_OFED does not support BlueField-3 devices.

- **[ConnectX-7] Multi-host LAG:** When using a multi-host deployment, each host is assigned unique ports and PFs and manages its own LAG.
- **Hardware GRO:** Hardware Generic Receive Offload (GRO) reduces CPU overhead and improves throughput by offloading the task of reassembling incoming network packets from the CPU.
- **Precision Time Measurement (PTM):** Precision Time Measurement (PTM) enhances clock synchronization accuracy between the NIC's hardware clock and the system's host clock. It addresses inaccuracies caused by asymmetric PCIe request and response times, which can affect PTP.
- **Transport RDMA Steering Domain:** Transport RDMA steering domain in the kernel allows users to create steering tables and rules using the new domain types FS\_FT\_RDMA\_TRANSPORT\_RX and FS\_FT\_RDMA\_TRANSPORT\_TX, with support for both DevX commands and the added functionality in `mlx5dv_create_flow_matcher()`.
- **DevX QPs Parallel Cleanup:** Reduces cleanup (kill) time of user-space applications that use DevX QPs, without requiring any changes on the user's end. When using a kernel that includes this patch, DevX QP destruction upon cleanup will be sent to the firmware in parallel, rather than serially as before, resulting in faster cleanup times. The more DevX QPs the application has, the more noticeable the improvement will be.
- **Flow Control for Solicited MADs:** Implements flow control for solicited MADs in the MAD layer to prevent loss due to insufficient resources at the receiver. Both the client and server function as receivers: the server handles incoming requests, while the client receives the responses.
- **New DOCA-Host OSEs:** RHEL 9.5, OpenEuler 22.03-SP4, OpenEuler 24.03-SP0, Oracle Linux 9.4, Oracle Linux 8.10, Azure Linux (Mariner) 3.0

## BlueField-3 Firmware Changes and New Features

- **PCIe Switch fwreset** - added support for a new synchronized flow, including a tool and driver, to perform a fwreset on setups with a PCIe switch configuration.
- **PTP** - Unified PTP is now supported across different VFs on the same PF
- **MADs** - added support for new MADs: `PortRecoveryPolicyConfig` and `PortRecoveryPolicyCounters`. During the PHY recovery process, the firmware core will indicate the `port_logical_state` as Active.
- **Multi-host MPFS LAG** - in multi-host DPU mode, all ports are assigned to host0, while only port 0 is assigned to the other hosts. A global LAG is supported to manage all hosts.
- **Data-Path Accelerator (DPA)** - the DPA hardware version is now exposed as a new capability, labeled "dpa\_platform\_version."
- **Dynamic Long Cables** - added the ability to set cable length as a parameter in the PFCC access register. The cable length is used in the calculation of RX lossless buffer parameters, including size, Xoff, and Xon thresholds.
- **Programmable Congestion Control (PCC): RTT, DSCP** - enables users to manually set the DSCP in the IP header of the RTT response using `mlxconfig`. Additionally, the PPCC register now supports querying the count of processed RTT request packets per port.
- **Block SMP Traffic** - Added a new NV config (`SM_DISABLE`, default 0) which, when enabled, blocks SMP traffic that does not originate from the SM

## BlueField-2 Firmware Changes and New Features


- **Dynamic Long Cables** - added the ability to set cable length as a parameter in the PFCC access register. The cable length is used in the calculation of RX lossless buffer parameters, including size, Xoff, and Xon thresholds.
- **Block SMP Traffic** - added a new NV config (`SM_DISABLE`, default 0) which, when enabled, blocks SMP traffic that does not originate from the SM.
- **PTP** - Unified PTP is now supported across different VFs on the same PF.

## API Changes

The following table presents the API changes which took place in this version of DOCA:

Library	Change Description
DOCA Comch	<p><b>Added</b> doca_comm_channel_2.h</p> <ul style="list-style-type: none"> <li>void doca_comm_channel_2_ep_set_comp_vector(struct doca_comm_channel_ep_t *ep, uint32_t comp_vector);</li> </ul>
DOCA Common	<p><b>Added</b> doca_ctx.h</p> <ul style="list-style-type: none"> <li>doca_error_t doca_ctx_set_completion_vector(struct doca_ctx *ctx, uint32_t comp_vector);</li> <li>doca_error_t doca_ctx_cap_get_num_completion_vectors(const struct doca_devinfo *devinfo, uint32_t *num_comp_vectors);</li> </ul> <p>doca_dev.h</p> <ul style="list-style-type: none"> <li>doca_error_t doca_devinfo_rep_get_iface_name(const struct doca_devinfo_rep *devinfo_rep, char *iface_name, uint32_t size);</li> </ul> <p>doca_mmap.h</p> <ul style="list-style-type: none"> <li>doca_error_t doca_mmap_enable_thread_safety(struct doca_mmap *mmap);</li> </ul> <p>doca_pe.h</p> <ul style="list-style-type: none"> <li>doca_error_t doca_pe_set_event_mode(struct doca_pe *pe, enum doca_pe_event_mode event_mode);</li> </ul> <p>doca_rdma_bridge.h</p> <ul style="list-style-type: none"> <li>doca_error_t doca_rdma_bridge_get_mmap_mkey_from_pd(const struct doca_mmap *mmap, struct ibv_pd *pd, uint32_t *mkey);</li> </ul> <p><b>Removed</b> doca_pe.h</p> <ul style="list-style-type: none"> <li>doca_error_t doca_pe_set_notification_affinity(struct doca_pe *pe, uint32_t core_id);</li> </ul>

Library	Change Description
DOCA DPA	<p><b>Changed</b> doca_dpa.h/doca_dpa_dev.h</p> <ul style="list-style-type: none"> <li>Removed DOCA_DPA_DEV_LOG_LEVEL_CRIT from enum doca_dpa_dev_log_level</li> </ul>
DOCA PCC	<p><b>Added</b> doca_pcc_np_dev.h</p> <ul style="list-style-type: none"> <li>uint8_t doca_pcc_np_dev_get_plane_num(const struct doca_pcc_np_dev_request_packet *input);</li> </ul>
DOCA RDMA	<p><b>Added</b> doca_rdma.h</p> <ul style="list-style-type: none"> <li>doca_error_t doca_rdma_set_rnr_retry_count(struct doca_rdma *rdma, uint8_t rnr_retry_count);</li> <li>doca_error_t doca_rdma_get_rnr_retry_count(const struct doca_rdma *rdma, uint8_t *rnr_retry_count);</li> </ul>

 **Info**

To view API changes in older DOCA versions, please refer to [their](#) specific release notes documentation.

# Customer Affecting Changes

## Changes in This Release

This section provides a list of changes that took place in the current version and break compatibility/interface or discontinue support for features or OS versions, etc.

Introduced in Version	Description	Customer Impact and Recommendation
DOCA-HOST 2.10.0 (Jan/2025)	When Dynamic Interrupt Moderation (DIM) is enabled, static coalescing parameters cannot be set, as their values will be overridden by the dynamic algorithm.	Disable Dynamic Interrupt Moderation (DIM) to set the coalescing parameters.
	ConnectX-4 adapter cards family is no longer supported.	N/A
	Removed support for the following OSes: <ul style="list-style-type: none"> <li>• RHEL8.0</li> <li>• RHEL8.1</li> <li>• RHEL8.3</li> <li>• RHEL8.5</li> <li>• RHEL8.7</li> <li>• RHEL9.1</li> <li>• RHEL9.3</li> <li>• OL 7.9</li> </ul>	N/A
DOCA-FLOW 2.10.0	<b>The following features are NOT supported in DOCA 2.10 release (The features would be supported in DOCA 3.0 April/25 release): ACL pipe, LPM pipe, CT pipe, Ordered_list pipe, external send queue (SQ), pipe resize</b>	Users can not use these DOCA Flow features. For applications that require this functionality, please use DOCA FLOW 2.9.1



Introduced in Version	Description	Customer Impact and Recommendation
	Removed the need to use a dummy_id in ipsec encryption action during pipe creation.	When creating a pipe with crypto action, UINT32_MAX will represent a changeable shared object. Otherwise, 0 < crypto_id < UINT32_MAX will indicate a constant. No use of dummy_id.
	The memory for modify field and encap actions needs to be allocated upfront per port	The doca_flow_port_cfg_set_actions_mem_size () function must be called to configure the size in case these actions are needed. The recommended initial mem size to provide can be calculated by: num_of_entries * DOCA_FLOW_MAX_ENTRY_ACTIONS_MEM_SIZE This can be tuned later to achieve better memory consumption
	Strict matching is no longer supported	The doca_flow_pipe_cfg_set_enable_strict_matching() function is deprecated. Setting enable_strict_matching to true in doca_flow_pipe_cfg_set_enable_strict_matching() is not supported anymore. Users need to do relaxed matching. More details and usage examples can be seen in section <a href="#">Relaxed Match</a> .
	The aging mechanism no longer relies on DPDK, thus the DPDK arg <i>service_core=&lt;num cores&gt;</i> , <i>svc_cycle_time=&lt;cycle time&gt;</i> are no longer supported	Users should use the following new API instead: doca_flow_port_cfg_set_service_threads_core() API to set the core number used for counter service doca_flow_port_cfg_set_service_threads_interval() API to set the counter service cycle interval

Introduced in Version	Description	Customer Impact and Recommendation
	The <code>doca_flow_parser_meta.random</code> field is changed to big-endian	Users should update it to big-endian. The recommended way is to use <code>DOCA_HTOBE16</code> for that
	Refactored RSS config API to adhere to other resource types	Users have to specify whether the RSS is shared or non-shared. Usage examples can be found in various samples, e.g. <code>flow_switch_rss_sample</code>
	Enumeration values were changed	Users need to recompile the application since the following enums values were changed: <ul style="list-style-type: none"> <li>1. <code>enum doca_flow_l2_meta</code></li> <li>2. <code>enum doca_flow_meter_color</code></li> </ul>
	DOCA FLOW Tune is at Alpha level, and switch visualization is not supported	N/A
	Before 2.10.0: During DOCA Flow initialization, the programs register a callback for processing important "life events" in the cycle of the entry. This callback is invoked upon "entry add" and "entry removal".	Starting 2.10.0: The callback is also invoked also upon "port stop" during the removal of pipe entries. Thus, user context should be still defined.
DOCA-ETH 2.10.0	<b>The following samples are NOT supported:</b> <b><code>eth_rxq_managed_mempool_receive</code></b> <b>and <code>eth_rxq_regular_receive</code></b>	Users can not use these DOCA samples. If these are needed, please use DOCA FLOW 2.9.1

## Changes Planned for Future Releases

This section provides a list of changes that will take place in a future version of the product and will break compatibility/interface or discontinue support for features or OS versions, etc.

Planned for Version	Description
N/A	N/A

## Changes in Earlier Releases

This section provides a list of changes that took place throughout the past two major releases that broke compatibility/interface or discontinued support for features or OS versions, etc.

### Info

For older changes, please consult the [archive](#) of DOCA documentation.

Introduced in Version	Description	Customer Impact and Recommendation
DOCA 2.9.0 (Oct/2024)	<p><b>DPA Outbox Blocking-Mode</b></p> <p>Due to a silicon issue, as of firmware version 28.43.2026, the DPA outbox is configured to operate in non-blocking mode, causing DPA outbox requests to complete immediately without waiting for actual completion. As a result, the DPA stack must poll a "busy" bit before initiating another DPA outbox operation.</p>	<p>Update the firmware version to 28.43.2026 or higher or update the BF-Bundle (containing this firmware) and DOCA-Host to 2.9.x or higher.</p> <p>This is mandatory for customers programming the DPA (e.g., DPA with DOCA PCC, or using NVIDIA turn-key apps which utilize the DPA (virtio-net/blk/fs, NVMe)).</p>
	<p><b>DPA Thread Context</b></p> <p>Due to internal-stack API changes, as of firmware v28.43.2026, DPA thread context is changed in the DPA. This affects the overlying DPA stack.</p> <p>As of firmware version 28.43.2026, internal-stack API changes have altered the DPA thread context, impacting the overlying DPA stack.</p>	

## Discontinued Features

List of features which are supported in previous generations of hardware devices.

N/A

---

# Declared Unsupported Features

The following DOCA Flow features are NOT supported in DOCA 2.10 release. They will be supported in DOCA 3.0 release (slated for April 2025).

## Note

Use DOCA FLOW 2.9.1 if any of these features are required.

- ACL pipe
- LPM pipe
- CT pipe
- Ordered list pipe
- External send queue (SQ)
- Pipe resize

**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.<br/><br/>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.<br/><br/>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.<br/><br/><br/><br/>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.<br/><br/><br/><br/>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.<br/><br/><br/><br/><b>Trademarks</b><br/><br/><br/>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.<br/>

© Copyright 2025, NVIDIA. PDF Generated on 05/05/2025