



# **NVIDIA MLNX\_OFED Documentation v24.04-0.6.6.0**

# Table of contents

<b>Release Notes</b>	7
General Support	11
Changes and New Features	26
Bug Fixes in This Version	28
Known Issues	30
<b>User Manual</b>	82
Introduction	82
Installation	94
Features Overview and Configuration	95
Programming	96
InfiniBand Fabric Utilities	99
Troubleshooting	110
Common Abbreviations and Related Documents	111
<b>Documentation History</b>	115
Release Notes History	115
User Manual Revision History	115
<b>Legal Notices and 3rd Party Licenses</b>	117

# List of Figures

Figure 0. Image2019 1 10 10 55 56 Version 1 Modificationdate  
1714752918737 Api V2

---

Figure 1. Procedure Heading Icon Version 1 Modificationdate  
1714752923413 Api V2

---

Figure 2. Procedure Heading Icon Version 1 Modificationdate  
1714752923413 Api V2

---

Figure 3. Worddave635fed9c99097774044df72a47e9130 Version 1  
Modificationdate 1714752931410 Api V2

---

Figure 4. Image2019 2 12 10 26 53 Version 1 Modificationdate  
1714752941043 Api V2

---

Figure 5. Procedure Heading Icon Version 1 Modificationdate  
1714752949833 Api V2

---

Figure 6. Procedure Heading Icon Version 1 Modificationdate  
1714752949833 Api V2

---

Figure 7. Procedure Heading Icon Version 1 Modificationdate  
1714752956733 Api V2

---

Figure 8. Image2022 7 26 10 29 13 Version 1 Modificationdate  
1714752972017 Api V2

---

Figure 9. Image2022 7 28 14 58 53 Version 1 Modificationdate  
1714752972490 Api V2

---

Figure 10. Image2022 7 28 15 0 34 Version 1 Modificationdate  
1714752974647 Api V2

---

Figure 11. Image2022 7 28 14 59 57 Version 1 Modificationdate  
1714752974020 Api V2

---

Figure 12. Image2022 7 28 15 1 11 Version 1 Modificationdate  
1714752975173 Api V2

---

Figure 13. Image2022 7 28 15 1 42 Version 1 Modificationdate  
1714752975570 Api V2

---

Figure 14. Image2022 7 28 15 2 14 Version 1 Modificationdate  
1714752975877 Api V2

---

Figure 15. Image2022 7 28 15 4 11 Version 1 Modificationdate  
1714752976543 Api V2

---

Figure 16. Image2022 7 28 15 6 40 Version 1 Modificationdate  
1714752977347 Api V2

---

Figure 17. Procedure Heading Icon Version 1 Modificationdate  
1714753022907 Api V2

---

Figure 18. Procedure Heading Icon Version 1 Modificationdate  
1714753022907 Api V2

---

Figure 19. Procedure Heading Icon Version 1 Modificationdate  
1714753022907 Api V2

---

Figure 20. Procedure Heading Icon Version 1 Modificationdate  
1714753022907 Api V2

---

Figure 21. Procedure Heading Icon Version 1 Modificationdate  
1714753022907 Api V2

---

Figure 22. Procedure Heading Icon Version 1 Modificationdate  
1714753038993 Api V2

---

Figure 23. Procedure Heading Icon Version 1 Modificationdate  
1714753038993 Api V2

---

Figure 24. Procedure Heading Icon Version 1 Modificationdate  
1714753038993 Api V2

---

Figure 25. Procedure Heading Icon Version 1 Modificationdate  
1714753038993 Api V2

---

Figure 26. Worddavb2ee67a7eb9aae5c536610e39a37dcc5 Version 1  
Modificationdate 1714753048077 Api V2

---

Figure 27. Worddav6931c32564b3b0c166f4a26788219144 Version 1  
Modificationdate 1714753046437 Api V2

---

Figure 28. Procedure Heading Icon Version 1 Modificationdate  
1714753046110 Api V2

---

Figure 29. Image2019 3 8 12 50 6 Version 1 Modificationdate  
1714753045200 Api V2

---

Figure 30. Procedure Heading Icon Version 1 Modificationdate  
1714753046110 Api V2

---

Figure 31. Procedure Heading Icon Version 1 Modificationdate  
1714753046110 Api V2

---

Figure 32. Procedure Heading Icon Version 1 Modificationdate  
1714753046110 Api V2

---

Figure 33. Procedure Heading Icon Version 1 Modificationdate  
1714753046110 Api V2

---

Figure 34. Procedure Heading Icon Version 1 Modificationdate  
1714753046110 Api V2

---

Figure 35. Procedure Heading Icon Version 1 Modificationdate  
1714753046110 Api V2

---

Figure 36. Procedure Heading Icon Version 1 Modificationdate  
1714753046110 Api V2

---

Figure 37. Procedure Heading Icon Version 1 Modificationdate  
1714753046110 Api V2

---

Figure 38. Procedure Heading Icon Version 1 Modificationdate  
1714753046110 Api V2

---

Figure 39. Procedure Heading Icon Version 1 Modificationdate  
1714753066933 Api V2

---

Figure 40. Worddav336f9b6791fd85e08c8e6897697cd75b Version 1  
Modificationdate 1714753068290 Api V2

---

Figure 41. Procedure Heading Icon Version 1 Modificationdate  
1714753072103 Api V2

---

Figure 42. Procedure Heading Icon Version 1 Modificationdate  
1714753072103 Api V2

---

Figure 43. Procedure Heading Icon Version 1 Modificationdate  
1714753072103 Api V2

---

Figure 44. Procedure Heading Icon Version 1 Modificationdate  
1714753084150 Api V2

---

# Overview

NVIDIA® OpenFabrics Enterprise Distribution for Linux (MLNX\_OFED) is a single Virtual Protocol Interconnect (VPI) software stack that operates across all NVIDIA network adapter solutions.

NVIDIA OFED (MLNX\_OFED) is an NVIDIA-tested and packaged version of OFED and supports two interconnect types using the same RDMA (remote DMA) and kernel bypass APIs called OFED verbs—InfiniBand and Ethernet. Up to 400Gb/s InfiniBand and RoCE (based on the RDMA over Converged Ethernet standard) over 10/25/40/50/100/200/400Gb/s are supported with OFED to enable OEMs and System Integrators to meet the needs of end users in the said markets.

Further information on this product can be found in the following MLNX\_OFED documents:

- [Release Notes](#)
- [User Manual](#)

## Software Download

Please visit [nvidia.com/en-us/networking](https://nvidia.com/en-us/networking) Products Software InfiniBand Drivers  
[NVIDIA MLNX\\_OFED](#)

## Document Revision History

For the list of changes made to the User Manual, refer to [User Manual Revision History](#).

For the list of changes made to the Release Notes, refer to [Release Notes History](#).

---

# Release Notes

## Release Notes Update History

Version	Date	Description
24.04-0.6.6.0	May 05, 2024	Initial release of this document version. This release introduces <a href="#">Changes and New Features</a> and <a href="#">Bug Fixes in This Version</a> .

### **Note**

As of MLNX\_OFED version v5.1-0.6.6.0, the following are no longer supported.

- ConnectX-3
- ConnectX-3 Pro
- Connect-IB
- RDMA experimental verbs libraries (mlnx\_lib)

To utilize the above devices/libraries, refer to version 4.9 long-term support (LTS).

Release Notes contain the following sections:

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



- [Known Issues](#)

## Supported NIC Speeds

The Linux Driver operates across all NVIDIA network adapter solutions supporting the following uplinks to servers:

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

Uplink/Adapter Card	Driver Name	Uplink Speed
ConnectX-4		<ul style="list-style-type: none"> <li>InfiniBand: SDR, QDR, FDR, FDR10, EDR</li> <li>Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 56GbE<sup>1</sup>, 100GbE</li> </ul>

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

## Package Contents

Package	Revision	Licenses
clusterkit	1.12.449-1.2401033	BSD
dpcp	1.1.46-1.2401033	BSD-3-Clause
hcoll	4.8.3227-1.2401033	Proprietary
ibarr	0.1.3-1.2401033	(GPL-2.0 WITH Linux-syscall-note) OR BSD-2-Clause
ibdump	6.0.0-1.2401033	BSD2+GPL2
ibsim	0.12-1.2401033	GPLv2 or BSD
ibutils2	2.1.1-0.1.MLNX20240128.g605c7811.2401033	Mellanox Confidential and Proprietary
iser	24.01-OFED.24.01.0.3.3.1	GPLv2
isert	24.01-OFED.24.01.0.3.3.1	GPLv2
kernel-mft	4.27.0-83	Dual BSD/GPL
knem	1.1.4.90mlnx3-OFED.23.10.0.2.1.1	BSD and GPLv2

Package	Revision	Licenses
libvma	9.8.51-1	GPLv2 or BSD
libxlio	3.21.2-1	GPLv2 or BSD
mlnx-en	24.01-0.3.3.0.g480e4c3	GPLv2
mlnx-ethtool	6.4-1.2401033	GPL
mlnx-iproute2	6.4.0-1.2401033	GPL
mlnx-nfsrdma	24.01-OFED.24.01.0.3.3.1	GPLv2
mlnx-nvme	24.01-OFED.24.01.0.3.3.1	GPLv2
mlnx- ofa_kernel	24.01-OFED.24.01.0.3.3.1	GPLv2
mlnx-tools	24.01-0.2401033	GPLv2 or BSD
mlx-steering- dump	1.0.0-0.2401033	GPLv2
mpitests	3.2.22-8f11314.2401033	BSD
mstflint	4.16.1-2.2401033	GPL/BSD
multiperf	3.0-3.0.2401033	BSD 3-Clause, GPL v2 or later
ofed-docs	24.01-OFED.24.01.0.3.3	GPL/BSD
ofed-scripts	24.01-OFED.24.01.0.3.3	GPL/BSD
openmpi	4.1.7a1-1.2401033	BSD
opensm	5.18.0.MLNX20240128.3f266a48- 0.1.2401033	GPLv2 or BSD
openvswitch	2.17.8-1.2401033	ASL 2.0 and LGPLv2+ and SISSL
perftest	24.01.0-0.38.gd185c9b.2401033	BSD 3-Clause, GPL v2 or later
rdma-core	2307mlnx47-1.2401033	GPLv2 or BSD
rshim	2.0.19-0.gbf7f1f2	GPLv2
sharp	3.6.0.MLNX20240128.e669b4e8- 1.2401033	Proprietary
sockperf	3.10-0.git5ebd327da983.2401033	BSD

Package	Revision	Licenses
srp	24.01-OFED.24.01.0.3.3.1	GPLv2
ucx	1.16.0-1.2401033	BSD
xpmem	2.7.3-1.2401033	GPLv2 and LGPLv2.1
xpmem-lib	2.7-0.2310055	LGPLv2.1

## General Support

### Supported Operating Systems

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Support Model	ASAP <sup>2</sup> OVS-Kernel SR-IOV	ASAP <sup>2</sup> OVS-DPDK SR-IOV	NFS over RDMA	NVMe	GPU Direct Storage (GDS)	UCX-CUDA Version
Alma 8.5	x86_64	4.18.0-348.12.2.EL8_5.X86_64	Community	✗	✗	✗	✗	✗	✗
Anolis OS 8.4	AArch64	4.18.0-348.2.1.AN8_4.AARCH64	Community	✗	✗	✗	✗	✗	✗
	x86_64	4.18.0-305.AN8.X86_64	Community	✗	✗	✗	✗	✗	✗
Anolis OS 8.6	AArch64	5.10.134+	Primary	✗	✗	✗	✗	✗	✗
	x86_64	5.10.134+	Primary	✗	✗	✗	✗	✗	✗
BCLINUX 21.10SP2	AArch64	4.19.90-2107.6.0.0098.oe	Primary	✗	✗	✗	✓	✗	✗

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Support Model	ASAP <sup>2</sup> OVS-Kernel SR-IOV	ASAP <sup>2</sup> OVS-DPDK SR-IOV	NFS over RDMA	NVMe	GPUDirect Storage (GDS)	UCX-CUDA Version
		1.bclinux.aarch64							
	x86_64	4.19.90-2107.6.0.0100.oe 1.bclinux.x86_64	Primary	✗	✗	✓	✓	✗	✗
CentOS Stream v8	AArch64	4.18.0-535.el8.aarch64	Community	✗	✗	✗	✗	✗	✗
	x86_64	4.18.0-535.el8.x86_64	Community	✗	✗	✗	✗	✗	✗
CentOS Stream v9	AArch64	5.14.0-407.el9.x86_64	Community	✗	✗	✗	✗	✗	✗
	x86_64	5.14.0-407.el9.aarch64	Community	✗	✗	✗	✗	✗	✗
CTYUNO S2.0	AArch64	4.19.90-2102.2.0.0062.ctl 2.aarch64	Primary	✗	✗	✗	✗	✗	✗
	x86_64	4.19.90-2102.2.0.0062.ctl 2.x86_64	Primary	✗	✗	✗	✗	✗	✗
CTYUNO S23.01	AArch64	5.10.0-136.12.0.86.ctl3. aarch64	Primary	✗	✗	✗	✗	✗	✗
	x86_64	5.10.0-136.12.0.86.ctl3.x 86_64	Primary	✗	✗	✗	✗	✗	✗

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Support Model	ASAP <sup>2</sup> OVS-Kernel SR-IOV	ASAP <sup>2</sup> OVS-DPDK SR-IOV	NFS over RDMA	NVMe	GPU Direct Storage (GDS)	UCX-CUDA Version
Debian1 0.8	AArch64	4.19.0-14-arm64	Primary	✗	✗	✗	✗	✗	✗
<b>Debian1 0.9</b>	x86_64	4.19.0-14-amd64	Primary	✗	✗	✗	✗	✗	✗
	x86_64	4.19.0-16-amd64	Primary	✓	✗	✗	✓	✗	✗
Debian1 0.13	AArch64	4.19.0-21-arm64	Primary	✗	✗	✗	✗	✗	✗
	x86_64	4.19.0-21-amd64	Primary	✗	✗	✗	✗	✗	✗
Debian1 1.3	AArch64	5.10.0-13-arm64	Primary	✓	✗	✗	✓	✗	✗
	x86_64	5.10.0-13-amd64	Primary	✓	✗	✗	✓	✗	✗
Debian1 2.1	AArch64	6.1.0-10-arm64	Primary	✓	✗	✗	✓	✗	✗
	x86_64	6.1.0-10-amd64	Primary	✓	✗	✗	✓	✗	✗
EulerOS2 .0sp9	AArch64	4.19.90-vhulk2006.2.0.h171.eulerosv2r9.arch64	Community	✗	✗	✗	✗	✗	✗
	x86_64	4.18.0-147.5.1.0.h269.eulerosv2r9.x86_64	Community	✗	✗	✗	✗	✗	✗
EulerOS2 .0sp10	AArch64	4.19.90-vhulk2110.1.0.h8	Community	✗	✗	✗	✗	✗	✗

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Support Model	ASAP <sup>2</sup> OVS-Kernel SR-IOV	ASAP <sup>2</sup> OVS-DPDK SR-IOV	NFS over RDMA	NVMe	GPU Direct Storage (GDS)	UCX-CUDA Version
		60.eulerosv2r10.aarch64	ty						
	x86_64	4.18.0-147.5.2.4.h694.eulerosv2r10.x86_64	Community	✗	✗	✗	✗	✗	✗
EulerOS2.0sp11	AArch64	5.10.0-60.18.0.50.h323.eulerosv2r11.aarch64	Primary	✗	✗	✗	✗	✗	✗
	x86_64	5.10.0-60.18.0.50.h323.eulerosv2r11.x86_64	Primary	✗	✗	✗	✗	✗	✗
EulerOS2.0sp12	AArch64	5.10.0-136.12.0.86.h1032.eulerosv2r12.aarch64	Primary	✗	✗	✗	✗	✗	✗
	x86_64	5.10.0-136.12.0.86.h1032.eulerosv2r12.x86_64	Primary	✗	✗	✗	✗	✗	✗
KYLIN10SP2	AArch64	4.19.90-24.4.v2101.ky10.aarch64	Primary	✗	✗	✗	✗	✗	✗
	x86_64	4.19.90-24.4.v2101.ky10.x86_64	Primary	✓	✗	✗	✗	✗	✗
KYLIN10SP3	AArch64	4.19.90-52.22.v2207.ky10	Primary	✓	✗	✗	✗	✗	✗

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Support Model	ASAP <sup>2</sup> OVS-Kernel SR-IOV	ASAP <sup>2</sup> OVS-DPDK SR-IOV	NFS over RDMA	NVMe	GPU Direct Storage (GDS)	UCX-CUDA Version
		.aarch64							
	x86_64	4.19.90-52.22.v2207.ky10.x86_64	Primary	✓	✗	✗	✗	✗	✗
<b>Mariner 2.0</b>	x86_64	5.15.148.2-2.cm2	Primary	✗	✗	✗	✗	✗	✗
Oracle Linux 7.9	x86_64	5.4.17-2011.6.2.el7uek.x86_64	Primary	✗	✗	✗	✗	✗	✗
Oracle Linux 8.4	x86_64	5.4.17-2102.201.3.el8uek.x86_64	Primary	✗	✗	✗	✗	✗	✗
Oracle Linux 8.6	x86_64	5.4.17-2136.307.3.1.el8uek.x86_64	Primary	✗	✗	✗	✗	✗	✗
Oracle Linux 8.7	x86_64	5.15.0-3.60.5.1.el8uek.x86_64	Primary	✗	✗	✗	✗	✗	✗
<b>Oracle Linux 8.8</b>	x86_64	5.15.0-101.103.2.1.el8uek.x86_64	Primary	✗	✗	✗	✗	✗	✗
Oracle Linux 9.0	x86_64	5.15.0-0.30.19.el9uek.x86_64	Primary	✗	✗	✗	✗	✗	✗
Oracle Linux 9.1	x86_64	5.15.0-3.60.5.1.el9uek.x86_64	Primary	✗	✗	✗	✗	✗	✗
<b>Oracle Linux 9.2</b>	x86_64	5.15.0-101.103.2.1.el9u	Primary	✗	✗	✗	✗	✗	✗



Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Support Model	ASAP <sup>2</sup> OVS-Kernel SR-IOV	ASAP <sup>2</sup> OVS-DPDK SR-IOV	NFS over RDMA	NVMe	GPU Direct Storage (GDS)	UCX-CUDA Version
		ek.x86_64							
OpenSUSE 15.3	AArch64	-	Community	✗	✗	✗	✗	✗	✗
	x86_64	5.3.18-150300.59.43-DEFAULT	Community	✗	✗	✗	✗	✗	✗
OPENULER20.03 SP1	AArch64	4.19.90-2012.4.0.0053.OE1.AARCH64	Community	✗	✗	✗	✗	✗	✗
	x86_64	4.19.90-2110.8.0.0119.OE1.X86_64	Community	✗	✗	✗	✗	✗	✗
OPENULER20.03 SP3	AArch64	4.19.90-2112.8.0.0131.oe1.aarch64	Primary	✗	✗	✗	✗	✗	✗
	x86_64	4.19.90-2112.8.0.0131.oe1.x86_64	Primary	✗	✗	✗	✗	✗	✗
OPENULER22.03	AArch64	5.10.0-60.18.0.50.oe2203.aarch64	Primary	✗	✗	✗	✗	✗	✗
	x86_64	5.10.0-60.18.0.50.oe2203.x86_64	Primary	✗	✗	✗	✗	✗	✗
Photon OS 3.0	x86_64	4.19.225-3.ph3	Community	✗	✗	✗	✗	✗	✗

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Support Model	ASAP <sup>2</sup> OVS-Kernel SR-IOV	ASAP <sup>2</sup> OVS-DPDK SR-IOV	NFS over RDMA	NVMe	GPU Direct Storage (GDS)	UCX-CUDA Version
RHEL/CentOS8.0	AArch64	4.18.0-80.el8.aarch64	Primary	✓	✓	✗	✗	✓	12.4
	x86_64	4.18.0-80.el8.x86_64	Primary	✓	✓	✗	✗	✓	12.4
RHEL/CentOS8.1	AArch64	4.18.0-147.el8.aarch64	Primary	✓	✓	✗	✗	✓	12.4
	x86_64	4.18.0-147.el8.x86_64	Primary	✓	✓	✗	✗	✓	12.4
RHEL/CentOS8.2	AArch64	4.18.0-193.el8.aarch64	Primary	✓	✓	✗	✓	✓	12.4
	x86_64	4.18.0-193.el8.x86_64	Primary	✓	✓	✓	✓	✓	12.4
RHEL/CentOS8.3	AArch64	4.18.0-240.el8.aarch64	Primary	✓	✓	✗	✗	✓	12.4
	x86_64	4.18.0-240.el8.x86_64	Primary	✓	✓	✗	✗	✓	12.4
RHEL/CentOS8.4	AArch64	4.18.0-305.el8.aarch64	Primary	✓	✓	✗	✓	✓	12.4
	x86_64	4.18.0-305.el8.x86_64	Primary	✓	✓	✓	✓	✓	12.4
RHEL/CentOS/Rocky8.5	AArch64	4.18.0-348.el8.aarch64	Primary	✓	✓	✗	✓	✓	12.4
	x86_64	4.18.0-348.el8.x86_64	Primary	✓	✓	✓	✓	✓	12.4
RHEL/Rocky8.6	AArch64	AArch64.18.0-372.41.1.el8_6.aarch64	Primary	✓	✓	✗	✓	✓	12.4

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Support Model	ASAP <sup>2</sup> OVS-Kernel SR-IOV	ASAP <sup>2</sup> OVS-DPDK SR-IOV	NFS over RDMA	NVMe	GPU Direct Storage (GDS)	UCX-CUDA Version
	x86_64	4.18.0-372.41.1.el8_6.x86_64	Primary	✓	✓	✓	✓	✓	12.4
RHEL/Rocky8.7	AArch64	4.18.0-425.14.1.el8_7.aarch64	Primary	✓	✗	✗	✓	✓	12.4
	x86_64	4.18.0-425.14.1.el8_7.x86_64	Primary	✓	✗	✓	✓	✓	12.4
RHEL/Rocky8.8	AArch64	4.18.0-477.10.1.el8_8.aarch64	Primary	✓	✗	✗	✓	✓	12.4
	x86_64	4.18.0-477.10.1.el8_8.x86_64	Primary	✓	✗	✓	✓	✓	12.4
RHEL/Rocky8.9	AArch64	4.18.0-513.5.1.el8_9.aarch64	Primary	✓	✗	✗	✓	✓	12.4
	x86_64	4.18.0-513.5.1.el8_9.x86_64	Primary	✓	✗	✓	✓	✓	12.4
RHEL/Rocky9.0	AArch64	5.14.0-70.46.1.el9_0.aarch64	Primary	✓	✗	✗	✓	✓	12.4
	x86_64	5.14.0-70.46.1.el9_0.x86_64	Primary	✓	✗	✓	✓	✓	12.4
RHEL/Rocky9.1	AArch64	5.14.0-162.19.1.el9_1.aarch64	Primary	✓	✗	✗	✓	✓	12.4

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Support Model	ASAP <sup>2</sup> OVS-Kernel SR-IOV	ASAP <sup>2</sup> OVS-DPDK SR-IOV	NFS over RDMA	NVMe	GPU Direct Storage (GDS)	UCX-CUDA Version
		rch64							
	x86_64	5.14.0-162.19.1.el9_1.x86_64	Primary	✓	✗	✓	✓	✓	12.4
RHEL/Rocky9.2	AArch64	5.14.0-284.11.1.el9_2.aarch64	Primary	✓	✗	✗	✓	✓	12.4
	x86_64	5.14.0-284.11.1.el9_2.x86_64	Primary	✓	✗	✓	✓	✓	12.4
RHEL/Rocky9.3	AArch64	5.14.0-362.8.1.el9_3.aarch64	Primary	✓	✗	✓	✓	✓	12.4
	x86_64	5.14.0-362.8.1.el9_3.x86_64	Primary	✓	✗	✓	✓	✓	12.4
RHEL/Rocky9.4	AArch64	5.14.0-427.13.1.el9_4.aarch64	Primary	✓	✗	✓	✓	✓	12.4
	x86_64	5.14.0-427.13.1.el9_4.x86_64	Primary	✓	✗	✓	✓	✓	12.4
SLES15S P2	AArch64	5.3.18-22-default	Primary	✓	✓	✗	✓	✗	✗
	x86_64	5.3.18-22-default	Primary	✓	✓	✓	✓	✗	✗
SLES15S P3	AArch64	5.3.18-57-default	Primary	✓	✗	✗	✓	✗	✗

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Support Model	ASAP <sup>2</sup> OVS-Kernel SR-IOV	ASAP <sup>2</sup> OVS-DPDK SR-IOV	NFS over RDMA	NVMe	GPU Direct Storage (GDS)	UCX-CUDA Version
	x86_64	5.3.18-57-default	Primary	✓	✗	✓	✓	✗	✗
SLES15S P4	AArch64	5.14.21-150400.22-default	Primary	✓	✗	✗	✓	✗	✗
	x86_64	5.14.21-150400.22-default	Primary	✓	✗	✓	✓	✗	✗
SLES15S P5	AArch64	5.14.21-150500.53-default	Primary	✓	✗	✗	✓	✗	✗
	x86_64	5.14.21-150500.53-default	Primary	✓	✗	✓	✓	✗	✗
tencentos3.3	AArch64	5.4.119-19.0009.39	Primary	✗	✗	✗	✓	✗	✗
	x86_64	5.4.119-19.0009.39	Primary	✗	✗	✗	✓	✗	✗
Ubuntu2 0.04	AArch64	5.4.0-26-generic	Primary	✓	✓	✗	✓	✓	12.4
	x86_64	5.4.0-26-generic	Primary	✓	✓	✓	✓	✓	12.4
Ubuntu2 2.04	AArch64	5.15.0-25-generic	Primary	✓	✓	✓	✓	✓	12.4
	x86_64	5.15.0-25-generic	Primary	✓	✓	✓	✓	✓	12.4
Ubuntu2 4.04	AArch64	6.6.0-14-generic	Primary	✓	✓	✓	✓	✓	12.4

Operating System	Architecture	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Support Model	ASAP <sup>2</sup> OVS-Kernel SR-IOV	ASAP <sup>2</sup> OVS-DPDK SR-IOV	NFS over RDMA	NVMe	GPU Direct Storage (GDS)	UCX-CUDA Version
	x86_64	6.6.0-14-generic	Primary	✓	✓	✓	✓	✓	12.4
UOS20.1060	AArch64	5.10.0-46.uel20.aarch64	Primary	✗	✗	✗	✗	✗	✗
	x86_64	5.10.0-46.uel20.x86_64	Primary	✗	✗	✗	✗	✗	✗
UOS20.1060a	AArch64	5.10.0-46.uelc20.aarch64	Primary	✗	✗	✗	✗	✗	✗
	x86_64	5.10.0-46.uelc20.x86_64	Primary	✗	✗	✗	✗	✗	✗
Citrix XenServer Host8.2	x86_64	4.19.0+1	Primary	✗	✗	✗	✗	✗	✗
Kernel 6.7	AArch64	6.7	Primary	✓	✗	✗	✓	✗	✗
	x86_64	6.7	Primary	✓	✗	✓	✓	✗	✗

**Note**

- 32 bit platforms are no longer supported in MLNX\_OFED
- For RPM-based distributions, to install OFED on a different kernel, create a new ISO image using `mlnx_add_kernel_support.sh` script (see the MLNX\_OFED User Manual for instructions)

- Upgrading MLNX\_OFED on a cluster requires upgrading all of its nodes to the newest version as well
- If using MLNX\_OFED 4.9 LTS with MLNX\_OFED 5.x with upstream verbs, MLNX\_OFED 4.9 must be installed with --upstream-libs flag so the verbs libraries match.
- A combination of 4.9 LTS default verbs and MOFED 5.x upstream verbs is not supported.
- All operating systems listed above are fully supported in Paravirtualized and SR-IOV environments with Linux KVM Hypervisor

## Upgrade/Downgrade Matrix

This section reflects which versions were tested and verified for upgrade and downgrade.

Target Version	Versions Verified for Upgrade/Downgrade	Release Type	Release Date
24.04-0.6.6.0	24.01-0.3.3.1 - DOCA-OFED Profile	GA	January 2024
	23.10-2.1.3.1 - MLNX_OFED	GA-LTS-Update	November 2023
	5.8-4.1.5.0	GA-LTS-Update	December 2023

## MLNX\_OFED Version Interoperability

This section reflects which versions were tested and verified for multi-version environments.

Target Version	Versions Verified for Upgrade/Downgrade	Release Type	Release Date
24.04-0.6.6.0	24.01-0.3.3.1 - DOCA-OFED Profile	GA	January 2024

Target Version	Versions Verified for Upgrade/Downgrade	Release Type	Release Date
	23.10-2.1.3.1 - MLNX_OFED	GA-LTS-Update	November 2023
	5.8-4.1.5.0	GA-LTS-Update	December 2023

## Supported NIC Firmware Versions

### Note

As of version 5.1, ConnectX-3, ConnectX-3 Pro or Connect-IB adapter cards are no longer supported. To work with a version that supports these adapter cards, please refer to version 4.9 long-term support (LTS).

This current version is tested with the following NVIDIA adapter card firmware versions:

Adapter Card	Bundled Firmware Version
BlueField®-2	24.41.1000
ConnectX-7	28.41.1000
ConnectX-6 Lx	26.41.1000
ConnectX-6 Dx	22.41.1000
ConnectX-6	20.41.1000
ConnectX-5/ConnectX-5 Ex	16.35.3502
BlueField	18.33.1048
ConnectX-4 Lx	14.32.1010
ConnectX-4	12.28.2006

For the official firmware versions, please see [https://www.nvidia.com/en-us/networking/Support Support Firmware Download](https://www.nvidia.com/en-us/networking/Support%20Support%20Firmware%20Download).



## Supported Non-Linux Virtual Machines

The following are the supported non-Linux Virtual Machines in this current version:

NIC	Windows Virtual Machine Type	Minimal WinOF Version	Protocol
ConnectX-7	All Windows server/client editions	MLNX_WinOF2 2.90	IPoIB, ETH
ConnectX-4 Lx / ConnectX-5 / ConnectX-6	All Windows server/client editions	MLNX_WinOF2 2.50	IPoIB, ETH

## Support in ASAP2—Accelerated Switch and Packet Processing®

ASAP <sup>2</sup> Requirements	<ul style="list-style-type: none"><li>• iproute &gt;= 4.12 (for tc support)</li><li>• Upstream Open vSwitch &gt;= 2.8 for CentOS 7.2 NVIDIA openvswitch</li></ul>
ASAP <sup>2</sup> -Supported Adapter Cards	<ul style="list-style-type: none"><li>• ConnectX-7</li><li>• ConnectX-6 Lx</li><li>• ConnectX-6 Dx</li><li>• ConnectX-5</li></ul>

## Lustre Versions Compatible with MLNX\_OFED

- Lustre 2.15.2
- Lustre 2.12.9

## Hardware and Software Requirements

- Linux operating system
- Administrator privileges on your machine(s)

- Disk Space: 1GB

For the OFED Distribution to compile on your machine, some software packages of your operating system (OS) distribution are required.

To install the additional packages, run the following commands per OS:

Operating System	Required Packages Installation Command
RHEL/Oracl e Linux	yum install perl pciutils python gcc-gfortran libxml2-python tcsh libnl.i686 libnl expat glib2 tcl libstdc++ bc tk gtk2 atk cairo numactl pkgconfig ethtool lsof
XenServer	yum install perl pciutils python libxml2-python libnl expat glib2 tcl bc libstdc++ tk pkgconfig ethtool
SLES 15	python ethtool libatk-1_0-0 python2-libxml2-python tcsh lib-stdc++6-devel-gcc7 libgtk-2_0-0 tcl libopenssl1_1 libnl3-200 make libcairo2 expat libmnl0 insserv-compat pciutils lsof lib-glib-2_0-0 pkg-config tk
Ubuntu/De bian	apt-get install perl dpkg autotools-dev autoconf libtool auto- make1.10 automake m4 dkms debhelper tcl tcl8.4 chrpath swig graphviz tcl-dev tcl8.4-dev tk-dev tk8.4-dev bison flex dpkg-dev curl libcurl4-gnutls-dev python-libxml2 libvirt-bin libvirt0 libnl-dev libglib2.0-dev libgfortran3 automake m4 pkg-config libnuma logrotate ethtool lsof

## Unsupported Functionalities/Features/NICs

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

- ConnectX-2 adapter card
- ConnectX-3 adapter card
- ConnectX-3 Pro adapter card
- Connect-IB adapter card
- Soft-RoCE
- 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

## Changes and New Features

The following are the new features and changes that were added in this version. The supported adapter cards are specified as follows:

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

For a list of features from previous versions see [Release Notes Change Log History](#) section.

Feature/Change	Description
24.04-0.6.6.0	
Live Migration	Added support for Live Migration. Live migration refers to the process of moving a guest virtual machine (VM) running on one physical host to another host without disrupting normal operations or causing any downtime or other adverse effects for the end user. For further information, see <a href="#">SR-IOV Live Migration</a> .
Linux Bridge Multicast Offload	Added support for multicast control path packet snooping (IGMP/MLD) and offload of MDB notifications in order to replicate multicast traffic to multiple destinations in the hardware.
Netdev Interface Maximum Channels	Increased the netdev maximum channels from 128 to 256 for systems with a high number of cores.
Page Management via Kernel's Page Pool	Enhanced performance by removing the internal Rx path from the internal page cache, thus allowing the driver to always use the kernel's page_pool.
XDP Enhancements	Enhanced performance by adding the following modifications in XDP: <ul style="list-style-type: none"> <li>• Added XDP multi-buffer support to the default RQ type (Striding RQ)</li> <li>• Added support for XDP multi-buffer redirect-in</li> <li>• Allowed non-linear single-segment frames in XDP TX MPWQE</li> </ul>
DC MRA Caps	Increased the maximum outstanding DC atomic reads to 32 to enable the driver to handle reads larger than 8K. Now <code>log_max_ra_{res/req}_dc</code> is exposed in DV API, enabling the user to see both RC and DC values.
SyncE userspace Support through Linux kernel DPLL Subsystem	Added support for Linux kernel DPLL subsystem as a mechanism for working with clock signals in NVIDIA's proprietary synchronous ethernet protocol daemon. This new mechanism enables the use of both VFs and SFs, and is supported starting from korg6.8.
Queued MADs Received	Limited the number of MADs received on the wire that are queued in the kernel as they are waiting for the user-space application to

Feature/Change	Description
	200k per umad file. By that, preventing malicious or faulty behavior of flooding a node with MADs and exhausting its memory.
Bridge Debuggability Extensions	Added the option to expose offloaded FDB state (flags, counters, etc.) via debugfs for debugging purposes. <pre>\$ cat mlx5/0000\:08\:00.0/esw/bridge/bridge1/fdb DEV MAC VLAN PACKETS BYTES LASTUSE FLAGS enp8s0f0_1 e4:0a:05:08:00:06 2 2 204 4295567112 0x0 enp8s0f0_0 e4:0a:05:08:00:03 2 3 278 4295567112 0x0</pre>
General	<a href="#">Bug fixes</a>

## Customer Affecting Changes

Customer Affecting Change	Description
Dynamic EQ	When an EQ is not created, the default affinity value is empty, therefore the "show" command output will show an empty affinity.

## API Changes in MLNX\_OFED

### MLNX\_OFED Verbs API Migration

As of MLNX\_OFED v5.0 release (Q1 of the year 2020), MLNX\_OFED Verbs API have migrated from the legacy version of user space verbs libraries (libibverbs, libmlx5, etc.) to the Upstream version rdma-core.

For the list of MLNX\_OFED verbs APIs that have been migrated, refer to [Migration to RDMA-Core document](#).

## Bug Fixes in This Version

Below are the bugs fixed in this version. For a list of fixes previous version, see [Bug Fixes History](#).

Internal Reference Number	Description
3798418	<b>Description:</b> Fixed an issue related to the comparison process between the SW steering and FW steering modes to avoid kernel crashes incidences.
	<b>Keywords:</b> SW steering, FW steering
	<b>Discovered in Release:</b> 23.10-1.1.9.0
	<b>Fixed in Release:</b> 24.04-0.6.6.0
3824463	<b>Description:</b> Fixed an issue that prevented the tc_wrap tool from properly working when VLAN is configured as the tool wrongly handled the library function return value.
	<b>Keywords:</b> tc_wrap tool, VLAN
	<b>Discovered in Release:</b> 23.10-1.1.9.0
	<b>Fixed in Release:</b> 24.04-0.6.6.0
3835880	<b>Description:</b> Fixed an IPsec issue that resulted in the disabling of Crypto offload.
	<b>Keywords:</b> IPsec, Crypto offload
	<b>Discovered in Release:</b> 23.10-1.1.9.0
	<b>Fixed in Release:</b> 24.04-0.6.6.0
3734180	<b>Description:</b> Fixed an issue that could have caused memory corruption when running XDP traffic.
	<b>Keywords:</b> Memory, XDP traffic
	<b>Discovered in Release:</b> 23.10-1.1.9.0
	<b>Fixed in Release:</b> 24.04-0.6.6.0
3818623	<b>Description:</b> Fixed an issue that caused IFA packets to be sent from the incorrect port on VF LAG.
	<b>Keywords:</b> IFA packets, VF LAG
	<b>Discovered in Release:</b> 23.10-1.1.9.0
	<b>Fixed in Release:</b> 24.04-0.6.6.0

Internal Reference Number	Description
3866587	<b>Description:</b> Added support for over KORG 6.7 read frequency_diff.
	<b>Keywords:</b> Operating System, frequency
	<b>Discovered in Release:</b> 23.10-1.1.9.0
	<b>Fixed in Release:</b> 24.04-0.6.6.0
3851731	<b>Description:</b> Fixed an issue that caused a deadlock in the flow of disabling the LAG when changing eswitch mode from switchdev to legacy when a LAG bond existed on the machine.
	<b>Keywords:</b> LAG
	<b>Discovered in Release:</b> 23.10-1.1.9.0
	<b>Fixed in Release:</b> 24.04-0.6.6.0
3546668	<b>Description:</b> In systems with 64k page size, applications opening a substantial amount of RDMA resources such as UARs, QPs, and CQs might encounter errors during the creation of these resources due to limitations in PCI BAR size.
	<b>Keywords:</b> PCI BAR size limitation
	<b>Discovered in Release:</b> 23.10-1.1.9.0
	<b>Fixed in Release:</b> 24.04-0.6.6.0
3764042	<b>Description:</b> Fixed an issue that resulted in PFC configuration being reset when updating the cable length via mlnx_qos.
	<b>Keywords:</b> PFC configuration, cable length
	<b>Discovered in Release:</b> 24.01-0.3.3.1
	<b>Fixed in Release:</b> 24.04-0.6.6.0

## Known Issues

The following is a list of general limitations and known issues of the current version of the release.

Internal Ref. Number	Issue
3916656	<p><b>Description:</b> If multicast join request is received during the Heavy Sweep, the InfiniBand Subnet Manager 5.19 may hang.</p>
	<p><b>Keywords:</b> Multicast, Subnet Manager</p>
	<p><b>Workaround:</b> Restart the Subnet Manager. For a comprehensive solution, upgrade to Subnet Manager version 5.19.1 which includes the fix from <a href="#">NVIDIA website</a> or contact NVIDIA Support.</p>
	<p><b>Discovered in Release:</b> 24.04-0.6.6.0</p>
3856101	<p><b>Description:</b> In Debian 12, using dhcpcd instead of dhclient to configure the network interface (using Networkmanager) will result in wrong network interface configuration.</p>
	<p><b>Keywords:</b> dhcpcd, dhclient, Debian 12, Networkmanager</p>
	<p><b>Workaround:</b> Use dhclient to configure the network interface.</p>
	<p><b>Discovered in Release:</b> 24.04-0.6.6.0</p>
3964215	<p><b>Description:</b> Driver might try to access privileged registers resulting in an error with syndrome.</p>
	<p><b>Keywords:</b> Unbind and bind the function or restart the driver.</p>
	<p><b>Workaround:</b> N/A</p>
	<p><b>Discovered in Release:</b> 24.04-0.6.6.0</p>
3640907	<p><b>Description:</b> When using a kernel version lower than v5.5, application termination on PCIe Gen5 servers could lead to kernel problems, such as IOMMU call traces, because of a lack of support in the AMD IOMMU kernel component.</p>
	<p><b>Keywords:</b> PCIe Gen5, IOMMU, Call Trace</p>
	<p><b>Workaround:</b> To resolve the issue either:</p> <ul style="list-style-type: none"> <li>• Add kernel parameter cmdline "iommu=pt"</li> </ul> <p>or</p> <ul style="list-style-type: none"> <li>• Use a kernel that includes this upstream patchset:</li> </ul>



Internal Ref. Number	Issue
	<p><a href="https://patchwork.kernel.org/project/linux-mediatek/cover/20190908165642.22253-1-murphyt7@tcd.ie/">https://patchwork.kernel.org/project/linux-mediatek/cover/20190908165642.22253-1-murphyt7@tcd.ie/</a></p> <p><b>Discovered in Release:</b> 24.04-0.6.6.0</p>
3004304	<p><b>Description:</b> Setting NVMe num_p2p_queues module parameter value to be greater than 0, may cause a harmless warning "<i>irq #XXX: nobody cared</i>" with Call Trace afterwards.</p> <p><b>Keywords:</b> NVMe, Call Trace, num_p2p_queues</p> <p><b>Workaround:</b> N/A</p> <p><b>Discovered in Release:</b> 24.01-0.3.3.1</p>
3735400	<p><b>Description:</b> The NVMF connect command does not work on IB setups when AR (Adaptive Routing) is enabled, since the PI (the Protection Information that is used by the NVMF) and AR are not supported simultaneously .</p> <p><b>Keywords:</b> NVMF connect, PI, Adaptive Routing</p> <p><b>Workaround:</b> Disable the AR at the opensm, or, alternatively, disable the PI at the nvme_rdma with a new module parameter.</p> <p><b>Discovered in Release:</b> 24.01-0.3.3.1</p>
3774149	<p><b>Description:</b> In some cases, there could be a race condition between RDMA_WRITE and shared memory write, leading to the MPI receiving invalid data with large messages or collective operations between ranks on the same node.</p> <p><b>Keywords:</b> Race condition, RDMA_WRITE, shared memory write</p> <p><b>Workaround:</b> Set UCX_RNDV_SCHEME=get_zcopy to force using RDMA_READ protocol.</p> <p><b>Discovered in Release:</b> 24.01-0.3.3.1</p>
3565433	<p><b>Description:</b> An error may occur when creating a DCI due to oversized WQEs. This is caused by a loose enforcement of the allowed max quantity of SGEs.</p> <p><b>Keywords:</b> DCI, SGEs</p>

Internal Ref. Number	Issue
	<b>Workaround:</b> N/A
	<b>Discovered in Release:</b> 24.01-0.3.3.1
3732632	<b>Description:</b> Geneve offload does not operate together with FLEX_PARSER.
	<b>Keywords:</b> Geneve offload, FLEX_PARSER
	<b>Workaround:</b> Make sure that the firmware is appropriately configured by verifying that the FLEX_PARSER_PROFILE_ENABLE mlxconfig flag is set to 0.
	<b>Discovered in Release:</b> 24.01-0.3.3.1
3644590	<b>Description:</b> When working in switchdev mode, the number of XFRM IN rules that can be added is limited to 2047.
	<b>Keywords:</b> switchdev mode, XFRM IN rules
	<b>Workaround:</b> N/A
	<b>Discovered in Release:</b> 24.01-0.3.3.1
3563584	<b>Description:</b> In case of a steering loop, the packet would loop indefinitely, causing a device hang.
	<b>Keywords:</b> Steering loop
	<b>Workaround:</b> Enable firmware infinite loop protection.
	<b>Discovered in Release:</b> 24.01-0.3.3.1

Internal Ref. Number	Issue
3678715	<b>Description:</b> When attempting to restart drivers using openibd service while the nvme_rdma module is loaded, the process may fail. This behavior is intentional, as unloading nvme_rdma during the driver restart can lead to connectivity issues in other applications within the setup.
	<b>Keywords:</b> openibd service, nvme_rdma module
	<b>Workaround:</b> Manually unload the nvme_rdma module before performing the driver restart. This can be achieved using the <code>modprobe -r</code>

Internal Ref. Number	Issue
	<p>nvme_rdma command.</p> <p><b>Discovered in Release:</b> 23.10-1.1.9.0</p>
3676223	<p><b>Description:</b> When using kernel version 4.12 or above, it is advised to run  echo 0 &gt; /sys/bus/pci/devices/0000\08\00.0/sriov_drivers_autoprobe to avoid VF probing</p> <p><b>Keywords:</b> VF probing</p> <p><b>Workaround:</b> N/A</p> <p><b>Discovered in Release:</b> 23.10-1.1.9.0</p>
3682658	<p><b>Description:</b> While using the RDMA-CM user application and the AF_IB parameter, the kernel uses only the first byte of the private data to set the CMA version. In such scenario, any user data written to this byte will be overwritten.</p> <p><b>Keywords:</b> RDMA-CM user application, AF_IB, private data</p> <p><b>Workaround:</b> Do not use AF_IB for application's private data.</p> <p><b>Discovered in Release:</b> 23.10-0.5.5.0</p>
3640082	<p><b>Description:</b> A potential null pointer dereference might occur due to a missing update in the PCI subsystem code when creating the maximum number of VFs.  All kernel versions lacking the following fix are impacted:  "PCI: Avoid enabling PCI atomics on VFs."</p> <p><b>Keywords:</b> Maximal VF number</p> <p><b>Workaround:</b> N/A</p> <p><b>Discovered in Release:</b> 23.10-0.5.5.0</p>
3653417	<p><b>Description:</b> When offloading IPsec policy rules while in legacy mode there are two options:</p> <ol style="list-style-type: none"> <li>1. Software steering - The software stack will handle the task, and no device offload will take place.</li> </ol>

Internal Ref. Number	Issue
	<p>2. Changing the steering mode to firmware steering will return unsupported.</p> <p><b>Keywords:</b> IPsec, legacy mode</p> <p><b>Workaround:</b> Perform a devlink reload after changing the steering mode.</p> <p><b>Discovered in Release:</b> 23.10-0.5.5.0</p>
3612274	<p><b>Description:</b> Currently, either IPsec offload or TC offload for a specific interface is allowed. The offloading TC rule to an interface will fail if an IPsec rule is already offloaded on it, and vice-versa.</p> <p><b>Keywords:</b> IPsec offload, TC offload</p> <p><b>Workaround:</b> N/A</p> <p><b>Discovered in Release:</b> 23.10-0.5.5.0</p>
3596126	<p><b>Description:</b> OVS mirroring of both egress and ingress together with modified TTL is not supported by Connectx-5 cards, and may cause packets checksum issues and errors in the dmesg command.</p> <p><b>Keywords:</b> OVS mirroring, Connectx-5</p> <p><b>Workaround:</b> N/A</p> <p><b>Discovered in Release:</b> 23.10-0.5.5.0</p>
3538463	<p><b>Description:</b> A Kernel ABI problem in Sles15SP4 may lead to issues during driver start. This impacts kernels starting from version 5.14.21-150400.24.11.1 up to version 5.14.21-150400.24.63.1 (July 2022 to May 2023), inclusive. For more information, see <a href="https://www.suse.com/support/kb/doc/?id=000021137">https://www.suse.com/support/kb/doc/?id=000021137</a>.</p> <p><b>Keywords:</b> Kernel ABI, Sles15SP4, driver start</p> <p><b>Workaround:</b> Upgrade to a kernel version newer than 5.14.21-150400.24.63.1 (May 2023).</p> <p><b>Discovered in Release:</b> 23.10-0.5.5.0</p>
3637252	<p><b>Description:</b> When running over REHL7.6 with excessive RDMA/RoCE workload, kernel warnings may be triggered.</p>

Internal Ref. Number	Issue
	<b>Keywords:</b> REHL7.6, RDMA, RoCE
	<b>Workaround:</b> N/A
	<b>Discovered in Release:</b> 23.10-0.5.5.0

Internal Ref. Number	Issue
3046655	<b>Description:</b> A package manager upgrade with zypper (on an SLES system) may prompt a question about vendor change from "Mellanox Technologies" to "OpenFabrics".
	<b>Keywords:</b> Installation, SLES
	<b>Workaround:</b> Either accept the prompted change, or add the /etc/zypp/vendors.d/mlnx_ofed file with the following content: [main] vendors = Mellanox,OpenFabrics
	<b>Discovered in Release:</b> 23.07-0.5.0.0
3392477	<b>Description:</b> The ConnectX-7 firmware embedded in this MLNX_OFED version cannot be burnt using the MLNX_OFED installer script.
	<b>Keywords:</b> ConnectX-7, MLNX_OFED installer script
	<b>Workaround:</b> Please download and install the dedicated firmware from the web <a href="https://network.nvidia.com/support/firmware/connectx7ib/">https://network.nvidia.com/support/firmware/connectx7ib/</a>
	<b>Discovered in Release:</b> 23.07-0.5.0.0
3532756	<b>Description:</b> The kernel may crash when restarting the driver while IP sec rules are configured.
	<b>Keywords:</b> IP sec
	<b>Workaround:</b> Flush the IP sec configuration before reloading the driver: ip xfrm state flush ip xfrm policy flush
	<b>Discovered in Release:</b> 23.07-0.5.0.0

Internal Ref. Number	Issue
3472979	<b>Description:</b> When a large number of virtual functions are present, the output of the "ip link show" command may be truncated.
	<b>Keywords:</b> virtual functions, ip link show
	<b>Workaround:</b> N/A
	<b>Discovered in Release:</b> 23.07-0.5.0.0
3413938	<b>Description:</b> When using the mlnx-sf script, creating and deleting an SF with the same ID number in a stressful manner may cause the setup to hang due to a race between the create and delete commands.
	<b>Keywords:</b> Hang; mlnx-sf
	<b>Workaround:</b> N/A
	<b>Discovered in Release:</b> 23.07-0.5.0.0
3461572	<b>Description:</b> Configuring Multiport Eswitch LAG mode can be performed only via devlink from this release onwards. The compat sysfs should not be used to configure mpesw LAG.
	<b>Keywords:</b> Multiport Eswitch, compat sysfs, mpesw LAG
	<b>Workaround:</b> N/A
	<b>Discovered in Release:</b> 23.07-0.5.0.0
3464337	<b>Description:</b> Simultaneously adding or removing TC rules while operating on kernel version 6.3 could potentially result in stability issues.
	<b>Keywords:</b> ASAP, rules, TC
	<b>Workaround:</b> Make sure the following fix is part of the kernel: <a href="https://lore.kernel.org/netdev/20230504181616.2834983-3-vladbu@nvidia.com/T/">https://lore.kernel.org/netdev/20230504181616.2834983-3-vladbu@nvidia.com/T/</a>
	<b>Discovered in Release:</b> 23.07-0.5.0.0
3469484	<b>Description:</b> Mirror and connection tracking (CT) offload actions are not supported simultaneously if the kernel version does not support hardware miss to TC actions. Thus, when performing a CT offload

Internal Ref. Number	Issue
	<p>test, the actual number of offloaded connections may be lower than expected.</p> <p><b>Keywords:</b> ASAP, CT offload</p> <p><b>Workaround:</b> Make sure to have the following offending commit in the tree:  net/sched: act_ct: offload UDP NEW connections  Make sure to to have <a href="https://www.spinics.net/lists/stable-commits/msg303536.html">https://www.spinics.net/lists/stable-commits/msg303536.html</a> in the kernel tree to fix this issue.</p> <p><b>Discovered in Release:</b> 23.07-0.5.0.0</p>
3473331	<p><b>Description:</b> When performing a CT offload test, the actual number of offloaded connections may be lower than expected.</p> <p><b>Keywords:</b> ASAP, CT offload</p> <p><b>Workaround:</b> The fix is external to the driver, make sure to have this commit in the tree:  offending commit: net/sched: act_ct: offload UDP NEW connections  Make sure you have:  <a href="https://www.spinics.net/lists/stable-commits/msg303536.html">https://www.spinics.net/lists/stable-commits/msg303536.html</a> in the kernel tree to fix this issue.</p> <p><b>Discovered in Release:</b> 23.07-0.5.0.0</p>
Internal Ref. Number	Issue
3360710	<p><b>Description:</b> Configuring PFC in parallel to buffer size and prio2buffer commands may lead to misalignment between firmware and software in regards to receiving buffer ownership.</p> <p><b>Keywords:</b> NetDev, PFC, Buffer Size, prio2buffer</p> <p><b>Workaround:</b> First, configure PFC on all ports, and then perform other needed QoS (i.e., buffer_size or prio2buffer) configurations accordingly.</p> <p><b>Discovered in Release:</b> 23.04-0.5.3.3</p>
3413879	<p><b>Description:</b> OpenSM may not be started automatically if chkconfig was not installed before OpenSM is installed. Note, however, that</p>

Internal Ref. Number	Issue
	<p>chkconfig will fail to install if the directory (rather than symbolic link to directory) /etc/init.d already exists (e.g., from a previous installation of MLNX_OFED).</p> <p><b>Keywords:</b> Installation, OpenSM, chkconfig</p> <p><b>Workaround:</b> Install chkconfig before installing MLNX_OFED. If installing it fails, make sure /etc/init.d does not exist at the time of installing it.</p> <p><b>Discovered in Release:</b> 23.04-0.5.3.3</p>
3424596	<p><b>Description:</b> On SLES 15.4, installing MLNX_OFED using a package repository (with zypper) may trigger an error message about missing dependency for '<a href="#">librte_eal.so.20.0()(64bit)</a>'. This is because the inbox package libdpdk-20_0 is being uninstalled as it is incompatible with the MLNX_OFED rdma-core packages.</p> <p><b>Keywords:</b> Installation, SLES 15.4</p> <p><b>Workaround:</b> Uninstall the relevant packages: 'zypper uninstall libdpdk-20_0' before installing MLNX_OFED. This will also remove the inbox openvswitch package.</p> <p><b>Discovered in Release:</b> 23.04-0.5.3.3</p>
3433416	<p><b>Description:</b> On systems that were installed with MLNX_OFED 5.9 or older and include a CUDA package (ucx-cuda / hcoll-cuda), an upgrade to MLNX_OFED 23.04 using the package manager ("yum") method will fail. This is because MLNX_OFED up to 5.9 is built with CUDA 11. MLNX_OFED 23.04 is built with CUDA 12 and those CUDA versions are incompatible.</p> <p><b>Keywords:</b> Installation, CUDA, yum</p> <p><b>Workaround:</b> Remove CUDA packages included with OFED (ucx-cuda, hcoll-cuda) before upgrading. This will allow to upgrade MLNX_OFED regardless of CUDA version installed. To install them later, CUDA 12 must be installed on the system.</p> <p><b>Discovered in Release:</b> 23.04-0.5.3.3</p>
3420831	<p><b>Description:</b> mlx-steering-dump is not supported on systems in which Python3 is not the default.</p>



Internal Ref. Number	Issue
	<b>Keywords:</b> mlx-steering-dump, Python3
	<b>Workaround:</b> N/A
	<b>Discovered in Release:</b> 23.04-0.5.3.3
3351989	<b>Description:</b> If the underlying persistent device name exceeds 15 characters in length, the operating system will not be able to perform renaming (i.e., the device name will remain "eth ").
	<b>Keywords:</b> Persistent Interface Names
	<b>Workaround:</b> Add the --copy-ifnames-udev flag to the OFED installation command. Note that this flag is only applicable if the persistent name provided by the kernel, without the 'np ' suffix, is 15 characters or fewer.
	<b>Discovered in Release:</b> 23.04-0.5.3.3

Internal Ref. Number	Issue
3324094	<b>Description:</b> When working in legacy rq (striding rq off), with large MTU > 3712, a 10-20% degradation in performance might be seen when running UDP stream with 64 bytes message size.
	<b>Keywords:</b> NetDev, MTU, UDP Stream
	<b>Workaround:</b> N/A
	<b>Discovered in Release:</b> 5.9-0.5.6.0
3313137	<b>Description:</b> Virtual Functions depend on Physical Functions for device access (e.g, firmware host PAGE management). In addition, VF may need to access safely the PF 'driver data' to use the command interface as in the VFIO usage to support live migration. While the PF is missing its driver, the VFs are completely unusable. As such, upon PF unload, the SR-IOV is disabled by the PF itself. This is the standard widely seen behavior in Linux drivers today.
	<b>Keywords:</b> Core, SR-IOV, VF, PF
	<b>Workaround:</b> N/A

Internal Ref. Number	Issue
	<b>Discovered in Release:</b> 5.9-0.5.6.0
3320947	<b>Description:</b> When the system is overloaded, there is a possibility that one hour will pass between the creation of DevLink port and its usage/assignment, due to some locking. This will trigger a trace starting with: "Type was not set for devlink port."
	<b>Keywords:</b> Core, DevLink, System Overload
	<b>Workaround:</b> N/A
	<b>Discovered in Release:</b> 5.9-0.5.6.0
3046222	<b>Description:</b> Installing OFED with Open vSwitch packages failed over Ubuntu22 OS with inbox Open vSwitch installed on it. Inbox Open vSwitch packages should be removed first.
	<b>Keywords:</b> Installation, Ubuntu22
	<b>Workaround:</b> Use --with-openswitch flag along with the installation command.
	<b>Discovered in Release:</b> 5.9-0.5.6.0
3262725	<b>Description:</b> Devlink reload while deleting namespace may cause a deadlock on kernels older than Linux-6.0.
	<b>Keywords:</b> Devlink, Namespace
	<b>Workaround:</b> N/A
	<b>Discovered in Release:</b> 5.9-0.5.6.0
3253255	<b>Description:</b> RHEL 7 does not include built-in support for Python3. There are two potential ways to install it, and both install a package with a different name: 1. EPEL for RHEL7: python36 2. RHEL extra repository Python3 support is needed for using Pyverbs and the Python support of Open vSwitch. MLNX_OFED assumes that on RHEL7.x, if using Python3, that python36 from EPEL is used (otherwise the optional Python3 support cannot be used).
	<b>Keywords:</b> RHEL7, Python3

Internal Ref. Number	Issue
	<p><b>Workaround:</b> To use Python3 support on RHEL7, install python36 from the RHEL7 EPEL repository.</p> <p><b>Discovered in Release:</b> 5.9-0.5.6.0</p>

Internal Ref. Number	Issue
3215514	<p><b>Description:</b> On EulerOS 2.0SP11, installation with the yum method may fail with an error that mlnx-iproute2 is missing a dependency on <a href="#">libdb-5.3.so()(64bit)</a>.</p> <p><b>Keywords:</b> Installation, EulerOS 2.0SP11, yum</p> <p><b>Workaround:</b> Install in advance the mlnx-iproute2 package with rpm and with the --nodeps option. For example: rpm -Uv --nodeps RPMS/mlnx-iproute2-5.19.0-1.58101.x86_64.rpm</p> <p><b>Discovered in Release:</b> 5.8- 1.0.1.1</p>
3191223	<p><b>Description:</b> In old kernels, /etc/init.d/openibd stop will fail because of an existing TC rule. Because mlx5_ib is already unloaded, mlx5_core and mlx5_ib will be in an inconsistent state.</p> <p><b>Keywords:</b> ASAP<sup>2</sup>, eSwitch, TC Rules</p> <p><b>Workaround:</b> Set eSwitch mode to legacy before enabling SR-IOV or reload mlx5_core to change eSwitch mode to legacy.</p> <p><b>Discovered in Release:</b> 5.8- 1.0.1.1</p>
3199628	<p><b>Description:</b> ping -6 -i &lt;interface name&gt; is broken in v5.18.</p> <p><b>Keywords:</b> NetDev, -i flag</p> <p><b>Workaround:</b> In all operating systems that are running Kernel 5.18 and below, remove the -i flag.</p> <p><b>Discovered in Release:</b> 5.8- 1.0.1.1</p>
3002932	<p><b>Description:</b> Jumbo MTU must be set on all uplinks (i.e., uplinks of *_sf and *_sf_r) at all times.</p> <p><b>Keywords:</b> NetDev, MTU, Uplink</p>

Internal Ref. Number	Issue
	<p><b>Workaround:</b> Configure jumbo MTU (9216) on all uplink-related interfaces.</p> <p><b>Discovered in Release:</b> 5.8- 1.0.1.1</p>
3130859	<p><b>Description:</b> The yum install method might be broken on installer regenerated with --add-kernel-support-build-only.</p> <p><b>Keywords:</b> Installation, yum</p> <p><b>Workaround:</b> Delete the original mlnx-ofed-all-5.* package and recreate the repository with: createrepo RPMS/</p> <p><b>Discovered in Release:</b> 5.8- 1.0.1.1</p>
3149387	<p><b>Description:</b> The package neohost-backend (included in MLNX_OFED) has a strict dependency on Python 2.7 and on the existence of /usr/bin/python. This dependency is because of a pre-installation test (which is a rather non-standard method) for /usr/bin/python will fail the installation if without Python 2.7. As a result, default installation of this on newer systems that do not have a default of Python 2 has been disabled. If there is an explicit request for this installation using the command-line option --with-neohost-backend, this sanity check will be overridden and there will be an attempt to install it regardless. On newer systems, there is likely to not be /usr/bin/python even if Python 2 is installed; as such its installation will fail.</p> <p><b>Keywords:</b> Installation, Python 2</p> <p><b>Workaround:</b> If neohost-backend is needed on a newer system, install Python 2 in advance and create the symbolic link /usr/bin/python -&gt; python2.</p> <p><b>Discovered in Release:</b> 5.8- 1.0.1.1</p>
3213777	<p><b>Description:</b> Oracle Enterprise Linux version 9.0 generates kernel module packages that have dependencies that are not provided by their own kernel RPM packages and thus are not installable.</p> <p><b>Keywords:</b> Installation, Oracle Enterprise Linux v9.0</p> <p><b>Workaround:</b> N/A</p>

Internal Ref. Number	Issue
	<b>Discovered in Release:</b> 5.8- 1.0.1.1
3229904	<b>Description:</b> Restart driver fails to load OFED modules after installing OFED on SLES15sp4 with errata kernel 5.14.21-150400.24.21-default.
	<b>Keywords:</b> Installation
	<b>Workaround:</b> Install OFED with --add-kernel-support flag.
	<b>Discovered in Release:</b> 5.8- 1.0.1.1
3189424	<b>Description:</b> VLAN naming is limited to 16 characters (like all other interface names). For names longer than 16 characters, the kernel generates its own interface name VLAN (VID).
	<b>Keywords:</b> Core, VLAN, Interface Name
	<b>Workaround:</b> Select a name which complies to the 16-characters limitation.
	<b>Discovered in Release:</b> 5.8- 1.0.1.1
3220855	<b>Description:</b> Creating external SFs on BF ARM when the host (x86) operating system does not support SFs may cause the host to crash.
	<b>Keywords:</b> Core, Scalable Functions
	<b>Workaround:</b> N/A
	<b>Discovered in Release:</b> 5.8- 1.0.1.1
3239291	<b>Description:</b> In some topologies, like logical partitions, mlxfwreset is not supported.
	<b>Keywords:</b> Core, mlxfwreset
	<b>Workaround:</b> N/A
	<b>Discovered in Release:</b> 5.8- 1.0.1.1

Internal Ref. Number	Issue
3114823	<b>Description:</b> The first attempt to create a new iSER connection fails with the following messages in dmesg:

Internal Ref. Number	Issue
	<p>iSCSI Login timeout on Network Portal &lt;iSER_Target_IP_ADDR&gt;:3260  isert: isert_get_login_rx: isert_conn 00000000e9239d52 interrupted before got login req</p> <p>After the error, the iSER Initiator connects to the Target successfully, but the memory allocated for the first connection is not freed correctly. As a result, the failed attempt also causes memory leakage.</p> <ul style="list-style-type: none"> <li>• <a href="https://kernel.org">kernel.org</a> Kernel 5.18</li> <li>• RHEL 9.0</li> <li>• RHEL 8.6</li> <li>• Ubuntu 22.04</li> <li>• SLES 15 SP4</li> </ul> <p>The error happens due to a bug in the scsi_transport_iscsi module, which is not a part of MLNX_OFED. As such, the issue cannot be fixed in MLNX_OFED. The bug is already fixed in kernel 5.19 by the commit f6eed15f3ea7 ("scsi: iscsi: Exclude zero from the endpoint ID range").</p> <p><b>Workaround:</b> Update the kernel if the above errors are experienced. If the issue is still reproduced after the kernel update, ask your distro support to apply the bug fix from the upstream kernel.</p> <p><b>Keywords:</b> iSER Initiator</p> <p><b>Discovered in Release:</b> 5.7-1.0.2.0</p>
3096911	<p><b>Description:</b> Installing chkconfig on Rhel9.0 with OFED using yum failed (chkconfig creates /etc/init.d sym link and OFED creates files in this directory, causing a conflict).</p> <p><b>Workaround:</b> Installing chkconfig before OFED.</p> <p><b>Keywords:</b> Installation</p> <p><b>Discovered in Release:</b> 5.7-1.0.2.0</p>
3100544	<p><b>Description:</b> On a RHEL9.x system, in some cases where inbox modules do not match for the drivers being build, rebuilding the drivers (--add-kernel-support) works, but fails to install the built package, with many errors such as: kernel(__rdma_block_iter_next) = 0x8e7528da is needed by mlnx-ofa_kernel-modules-5.6-OFED.5.6.2.0.9.1.kver.5.14.0_70.13.1.el9_0.aarch64.aarch64</p>

Internal Ref. Number	Issue
	<p>This was caused by a bug in the scripts that creates the Requires and Provides headers that is confused by dependencies between different modules of the same external package.</p> <p><b>Workaround:</b> dnf install kernel-modules- # in case it is not the newest.</p> <p><b>Keywords:</b> Installation, RHEL9.x</p> <p><b>Discovered in Release:</b> 5.7-1.0.2.0</p>
3132158	<p><b>Description:</b> Building rdma-core package on Rocky 8.6 OS caused failure in OFED build.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Installation</p> <p><b>Discovered in Release:</b> 5.7-1.0.2.0</p>
3137440	<p><b>Description:</b> Python package is missing, need to install it manually.</p> <p><b>Workaround:</b> Install Python before starting the build.</p> <p><b>Keywords:</b> Installation, Python</p> <p><b>Discovered in Release:</b> 5.7-1.0.2.0</p>
3141506	<p><b>Description:</b> kernel-macros package does not support building with KMP enabled. KMP needs to be disabled.</p> <p><b>Workaround:</b> Build and install MOFED with KMP disabled (without --kmp flag).</p> <p><b>Keywords:</b> Installation</p> <p><b>Discovered in Release:</b> 5.7-1.0.2.0</p>
3141506	<p><b>Description:</b> kernel-macros package does not support building with KMP enabled. KMP needs to be disabled.</p> <p><b>Workaround:</b> Build and install MOFED with KMP disabled (without --kmp flag).</p> <p><b>Keywords:</b> Installation</p> <p><b>Discovered in Release:</b> 5.7-1.0.2.0</p>
3129627	<p><b>Description:</b> Kernel module packaging is not supported in CtyunOS.</p>

Internal Ref. Number	Issue
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Installation
	<b>Discovered in Release:</b> 5.7-1.0.2.0



Internal Ref. Number	Issue
2971708	<p><b>Description:</b> For OSs in which Devlink supports setting roce-enable/disable, both sysfs roce_enable show and sysfs roce_enable set are disabled, and the RoCE state must be managed exclusively via Devlink.</p> <p>The sysfs interface for roce-enable/disable will be removed entirely for these OSs in a future release.</p> <p>To determine if Devlink can be used to enable or disable RoCE, execute the following console command after starting OFED:</p> <pre>devlink dev param show   grep roce</pre> <p>Devlink supports roce enable/disable if the following line is reflected in the output:</p> <pre>name enable_roce type generic</pre> <p>For OSs which do not allow enabling/disabling RoCE via Devlink, the sysfs interface behaves as in the previous 2 releases:</p> <ol style="list-style-type: none"> <li>1. For OSs which have Devlink reload, but do not allow setting RoCE state via Devlink: sysfs roce_enable show works, as does sysfs roce_enable set, but Devlink reload must be performed after setting the RoCE state via sysfs in order to activate the desired roce state.</li> <li>2. For OSs which do not have Devlink reload, RoCE state is managed only by the sysfs interface. 'show' displays the RoCE state and 'set' sets the state and activates it.</li> </ol> <p>To determine if Devlink dev reload is supported, execute the following console command (using the bash shell):</p> <pre>devlink dev help 2&gt;&amp;1   grep reload</pre> <p>Reload is supported if the output is:</p> <pre>devlink dev reload DEV [ netns { PID   NAME   ID } ]</pre>
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Enabling/Disabling RoCE

Internal Ref. Number	Issue
	<b>Discovered in Release:</b> 5.7-1.0.2.0

Internal Ref. Number	Issue
2998194	<b>Description:</b> On some systems with many (e.g., 64) virtual functions (VFs) attached to a ConnectX interface, 'ip link' may give an error message: "Error: Buffer too small for object." This applies to both IP commands: the inbox iproute package in RHEL8.x and the mlnx-iproute2 package from MLNX_OFED. This is known to work well and not give an error in RHEL7.x kernel regardless of what user-space package is used (including user-space from RHEL8.x).
	<b>Workaround:</b> N/A
	<b>Keywords:</b> NetDev, RHEL, Virtual Functions
	<b>Discovered in Release:</b> 5.6-1.0.3.3
3045436	<b>Description:</b> Rebooting the host while the Arm is down may block the shutdown flow till the Arm is up.
	<b>Workaround:</b> Restart the driver on the host side before reboot.
	<b>Keywords:</b> Reboot, Arm
	<b>Discovered in Release:</b> 5.6-1.0.3.3
3040350	<b>Description:</b> <ol style="list-style-type: none"> <li>1. When offload is enabled, removing a physical port from ovs-dpdk bridge requires restarting OVS service. Not doing so will result in wrong configuration of datapath rules.</li> <li>2. When offload is enabled, the physical port must be attached to a bridge.</li> </ol>
	<b>Workaround:</b>

Internal Ref. Number	Issue
	<p>1. When removing a physical port from an ovs-dpdk bridge while offload is enabled, need to restart openvswitch after reattaching it. 2. Attach physical port to a bridge according to the desired topology.</p> <p><b>Keywords:</b> OVS-DPDK, Bridge, Offload</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
2973726	<p><b>Description:</b> dec_ttl only work with ConnectX-6. It does not work with ConnectX-5.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> OVS-DPDK, dec_ttl</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
2946873	<p><b>Description:</b> Moving to switchdev mode while deleting namespace may cause a deadlock.</p> <p><b>Workaround:</b> Unload mlx5_ib module before moving to Switchdev mode.</p> <p><b>Keywords:</b> ASAP<sup>2</sup>, Switchdev, Namespace</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
2811957	<p><b>Description:</b> If a system is run from a network boot and is connected to the network storage through an NVIDIA ConnectX card, unloading the mlx5_core driver (such as running '/etc/init.d/openibd restart') will render the system unusable and should therefore be avoided.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Installation, mlx5_core</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
2979243	<p><b>Description:</b> The kernel in CentOS 7.6alt (for non-x86 architectures) is different than that of RHEL 7.6alt. Some of the MLNX_OFED kernel modules that were built for the RHEL7.6alt kernel will not load on a system with Centos7.6alt kernel. If you want to install MLNX_OFED on such a system, you should use ./mlnxofedinstall --add-kernelsupport to rebuild the kernel modules for the Centos kernel.</p>

Internal Ref. Number	Issue
	<p><b>Workaround:</b> Use add-kernel-support.</p> <p><b>Keywords:</b> Installation,CentOS</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
3011440	<p><b>Description:</b> In Debian 11.2, Ubuntu 21.10, and Ubuntu 22.04, attempting to install an "exact" type of metapackage (such as mlnx-ofed-all-exact or mlnx-ofed-basic-exact) may fail with an error regarding the version of mstflint.</p> <p><b>Workaround:</b> Install also mstflint of the exact same version (e.g., apt install mlnx-ofed-all-exact mstflint=4.16.0-1.56xxxx).</p> <p><b>Keywords:</b> Installation,Debian, Ubuntu, MST</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
3024520	<p><b>Description:</b> The option --copy-ifnames-udev copy some files under /etc (/etc/udev/rules.d/82-net-setup-link.rules and /etc/infiniband/vf-net-link-name.sh) that are never removed--not in the case this option is not given and not upon uninstallation. Those scripts are merely examples. They are files under /etc to be maintained by the user.</p> <p><b>Workaround:</b> Remove the files, if needed.</p> <p><b>Keywords:</b> Installation</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
3046601	<p><b>Description:</b> When rebuilding the kernel modules (--add-kernel-support) for some kernel versions (specifically mainline 4.14) do not unset LDFLAGS properly. Rebuilding xpmem in such a case may fail with the error such as "unrecognized option '-Wl,-z,relro'" in the xpmem build log.</p> <p><b>Workaround:</b> Either disable building xpmem by adding --without-xpmem to the command line, or edit the kernel Makefile to make it unset LDFLAGS:</p> <pre data-bbox="391 1713 1463 1801">sed -i -e '/^export ARCH/iLDFLAGS :=' /lib/modules/\$(uname -r)/Makefile</pre> <p><b>Note:</b> The Makefile may be located elsewhere, such as the top-level directory of the kernel source directory.</p>

Internal Ref. Number	Issue
	<p><b>Keywords:</b> Installation, SLES</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
3048411	<p><b>Description:</b> After installing OFED with rebuilt kernel modules, error messages indicating that the kernel module mlx5_ib failed to load (e.g. "mlx5_ib: Unknown symbol . . .") appear. These messages could be safely ignored because the module eventually loads.</p> <p><b>Workaround:</b> Run the command 'dracut -f' to update the initramfs.</p> <p><b>Keywords:</b> Installation</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
3048444	<p><b>Description:</b> OFED installation failed using yum for --add-kernel-support option (building packages without KMP enabled) if libfabric package is installed.</p> <p><b>Workaround:</b> Remove libfabric package before OFED installation or use installation script.</p> <p><b>Keywords:</b> Installation, RHEL 8.5</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
3015210	<p><b>Description:</b> OVS topology where the tunnel device is over a VF and the VF representor is connected to a bond is not supported.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> ASAP<sup>2</sup>, Tunnel Over VF, LAG, Connection Tracking</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
3028300	<p><b>Description:</b> OVS metering is not support over kernel 5.17.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> ASAP<sup>2</sup>, OVS, Meter, Kernel 5.17</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
3044255	<p><b>Description:</b> Destroying mlxdevm group while SF is attached to it is not supported.</p> <p><b>Workaround:</b> N/A</p>

Internal Ref. Number	Issue
	<p><b>Keywords:</b> ASAP<sup>2</sup>, mlxdevm, QoS, Group, Scalable Functions, ConnectX-6 Dx</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
2900346	<p><b>Description:</b> On Ubuntu OS, configuring different IP addresses with different subnets to both ports 0 and 1 is currently not supported. When trying to ping from port 0 on one BlueField-2 card to port 0 on the other BlueField-2 card, then both port 0 and port 1 on the receiving side send a reply to the ARP request (a.k.a, ARP flux).</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> BlueField-2, Ubuntu, ARP Flux</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
3046456	<p><b>Description:</b> Switching between SwitchDev mode and legacy mode quickly on BlueField-2 can prevent the driver from loading successfully and breaks its health recovery.</p> <p><b>Workaround:</b> Pause 60 seconds between state-altering commands to guarantee the driver health recovery is completed successfully.</p> <p><b>Keywords:</b> ASAP<sup>2</sup>, Health Recovery</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
2934149	<p><b>Description:</b> Adding vDPA ports over ConnectX-5 devices in ovs-dpdk is not supported and will cause a crash.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> OVS-DPDK, ConnectX-5</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>
2934833	<p><b>Description:</b> Running I/O traffic and toggling both physical ports status (UP/DOWN) in a stressful manner on the receiving-end machine may cause traffic loss.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> RDMA, Port Toggle</p> <p><b>Discovered in Release:</b> 5.6-1.0.3.3</p>

Internal Ref. Number	Issue
2901514	<b>Description:</b> Relaxed Ordering is not working properly on Virtual Functions.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Relaxed Ordering, VF
	<b>Discovered in Release:</b> 5.6-1.0.3.3

Internal Ref. Number	Issue
2870299	<b>Description:</b> Managing SFs is possible using the iproute2 with mlxdevm tool only.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Scalable Functions
	<b>Discovered in Release:</b> 5.5-1.0.3.2
2869722	<b>Description:</b> OFED packages were built with DKMS disabled since building OFED with DKMS failed due to a problem in the DKMS package on UOS. --dkms flag should not be used.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Installation, DKMS
	<b>Discovered in Release:</b> 5.5-1.0.3.2
2870367	<b>Description:</b> On UOS, IPoIB PKEY may require manual bring up after driver restart.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Installation, IPoIB, PKEY

Internal Ref. Number	Issue
	<b>Discovered in Release:</b> 5.5-1.0.3.2
2836032	<p><b>Description:</b> When using Software steering mlx5dv_dr API to create rules containing encapsulation actions in MLNX_OFED v5.5-1.x.x.x, upgrade firmware to the latest version. Otherwise, the maximum number of encapsulation actions that can be created will be limited to only 16K, and degradation for the rule insertion rate is expected compared to MLNX_OFED v5.4-.x.x.x.x.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Software Steering</p> <p><b>Discovered in Release:</b> 5.5-1.0.3.2</p>
2851639	<p><b>Description:</b> Enabling ARFS in legacy mode and then moving to switchdev mode is not supported and may cause unwanted behavior.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> NetDev, ARFS</p> <p><b>Discovered in Release:</b> 5.5-1.0.3.2</p>
2851639	<p><b>Description:</b> nvme and iser are not enabled on UOS ARM, because of missing UOS kernel support.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> nvme, iser, UOS ARM</p> <p><b>Discovered in Release:</b> 5.5-1.0.3.2</p>
2860855	<p><b>Description:</b> Building OFED on RHEL 8.4 with kmp disabled and then installing with yum fails due to some conflicting packages.</p> <p><b>Workaround:</b> Remove libfabric and librpmem packages before OFED installation, or add --allowerasing option to the installation command.</p> <p><b>Keywords:</b> Installation, RHEL 8.4, kmp, yum</p> <p><b>Discovered in Release:</b> 5.5-1.0.3.2</p>
2865983	<p><b>Description:</b> OFED packages were built with kmp disabled. Building with kmp enabled fails due to missing packages.</p> <p><b>Workaround:</b> N/A</p>



Internal Ref. Number	Issue
	<b>Keywords:</b> Installation, kmp
	<b>Discovered in Release:</b> 5.5-1.0.3.2

Internal Ref. Number	Issue
2658644	<b>Description:</b> Only match on lower 32 bit of ct_label is supported.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> ASAP <sup>2</sup> , Connection Tracking
	<b>Discovered in Release:</b> 5.4-1.0.3.0
2706345	<b>Description:</b> Number of RQ and TIR allocation in the driver depends on total number of MSI-X vectors allocated. Total number of TIRs supported by device is 16K range. Each representor needs number of CPUs worth TIRs, upto maximum of 128.
	<b>Workaround:</b> To use large number of VFs, set PF_NUM_PF_MSIX to a smaller value of around 32.
	<b>Keywords:</b> ASAP <sup>2</sup> ,VF, PF_NUM_PF_MSIX
	<b>Discovered in Release:</b> 5.4-1.0.3.0
2836997	<b>Description:</b> An automatic test that checks a flow meter rate fluctuation stays within a fixed threshold (e.g., 10%) may fail because meter precision is dependent on multiple factors (i.e., rate and burst values and shape of the traffic). To pick the best configuration parameters for a flow meter, perform a couple of test measurements using different values of burst size against expected traffic workload and average the results over an extended period of time (tens of minutes).
	<b>Workaround:</b> N/A
	<b>Keywords:</b> ASAP <sup>2</sup> ,Meter Threshold
	<b>Discovered in Release:</b> 5.4-1.0.3.0
2863456	<b>Description:</b> SA limit by packet count (hard and soft) are supported only on traffic originated from the ECPF. Trying to configure them on VF

Internal Ref. Number	Issue
	<p>traffic will remove the SA when hard limit is hit, however traffic could still pass as plain text due to the tunnel offload that is used in such configuration.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> ASAP<sup>2</sup>, IPsec Full Offload</p> <p><b>Discovered in Release:</b> 5.4-0.5.1.1</p>
2657392	<p><b>Description:</b> OFED installation caused CIFS to break in RHEL 8.4 and above. A dummy module was added so that CIFS will be disabled after OFED installation in RHEL 8.4 and above.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Installation, RHEL, CIFS</p> <p><b>Discovered in Release:</b> 5.4-0.5.1.1</p>
2800993	<p><b>Description:</b> OpenMPI does not support running across different operating systems and/or CPU architectures.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> OpenMPI</p>
2399503	<p><b>Description:</b> Open vSwitch is not supported on the latest operating systems containing only Python3 support.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Python, Open vSwitch</p>
2657392	<p><b>Description:</b> OFED installation caused CIFS to break in RHEL8.4. A dummy module was added so that CIFS will be disabled after OFED installation in RHEL8.4.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Installation, RHEL8.4, CIFS</p> <p><b>Discovered in Release:</b> 5.4-0.5.1.1</p>
2782406	<p><b>Description:</b> Running yum update will upgrade kylin-release to a higher version. The version of this package is used for kylin10sp2 detection so</p>

Internal Ref. Number	Issue
	<p>the script will detect kylin 10 instead of kylin10sp2 and use its repository by mistake.</p> <p><b>Workaround:</b> Because there are no special cases for kylin10sp2, the repository that was detected with adding --add-kernel-support to the installation command can be used.</p> <p><b>Keywords:</b> Upgrade, kylin</p> <p><b>Discovered in Release:</b> 5.4-3.0.3.0</p>
2755632	<p><b>Description:</b> On dual port cards with SR-IOV, when one port link is configured to InfiniBand and the other port link is configured to Ethernet, the Ethernet port will not be able to support VST and QinQ.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> SR-IOV, VST, QinQ</p> <p><b>Discovered in Release:</b> 5.4-3.0.3.0</p>
2780436	<p><b>Description:</b> Non-default MTU (&gt;1500) is not supported with IPsec crypto offload and may cause packet drops.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> IPsec, Crypto Offload, MTU</p> <p><b>Discovered in Release:</b> 5.4-3.0.3.0</p>
2726021	<p><b>Description:</b> Building packages on openEuler with kmp enabled requires kernel-rpm-macros package installed. kernel-rpm-macros-30-13.oe1 does not support -p option and kernel-rpm-macros-30-18.oe1 should be installed instead.</p> <p>On kylin OS, the version of kernel-rpm-macros package does not support -p option needed to support kmp, so it will stay disabled.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Installation, openEuler</p> <p><b>Discovered in Release:</b> 5.4-3.0.3.0</p>

Internal Ref. Number	Issue
2750653	<b>Description:</b> Running fragmented traffic in RHEL 8.3 (4.18.0-240.el8.x86_64) may cause call trace in build_skb.
	<b>Workaround:</b> Update to RHEL 8.3 z-stream 4.18.0-240.22.1.el8_3.x86_64.
	<b>Keywords:</b> RHEL 8.3, Kernel Panic, Call Trace, fr
	<b>Discovered in Release:</b> 5.4-1.0.3.0
2629375	<b>Description:</b> Matching on CT label is only supported when matching on lower 32 bits. Full match on all 128 bits of CT label is not supported.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> ASAP <sup>2</sup> , Connection Tracking, Label
	<b>Discovered in Release:</b> 5.4-1.0.3.0
2707997	<b>Description:</b> Installation in the package manager mode under SLES 15.x may require user-intervention if the original libibverbs is installed.
	<b>Workaround:</b> zypper install --force-resolution mlnx-ofed-all
	<b>Keywords:</b> Installation, libibverbs
	<b>Discovered in Release:</b> 5.4-1.0.3.0
2708531	<b>Description:</b> Installation in the package manager mode under SLES 15.x may require user-intervention if the original libopenvswitch is installed.
	<b>Workaround:</b> zypper install --force-resolution mlnx-ofed-all
	<b>Keywords:</b> Installation
	<b>Discovered in Release:</b> 5.4-1.0.3.0
2703043	<b>Description:</b> Congested TCP lock for kTLS TX device offload traffic compromises the performance.
	<b>Workaround:</b> Disable TCP selective acknowledgement: echo 0 > /proc/sys/net/ipv4/tcp_sack
	<b>Keywords:</b> kTLS TX
	<b>Discovered in Release:</b> 5.4-1.0.3.0

Internal Ref. Number	Issue
2676405	<b>Description:</b> If the package interface-rename is active (on XenServer, for example), the interface renaming by the OFED will not be done to eliminate conflicts.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Interface Renaming
	<b>Discovered in Release:</b> 5.4-1.0.3.0
2687943	<b>Description:</b> Offload of rules which redirect from VF on one PF to VF on second PF is not supported on socket-direct devices.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> ASAP <sup>2</sup> , Socket-Direct
	<b>Discovered in Release:</b> 5.4-1.0.3.0
2678672	<b>Description:</b> When disabling switchdev mode, the qdisc in tunnel device cannot be destroyed and mlx5e_stats_flower() is still called by OVS resulting in NULL pointer panic and memory leak.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> SwitchDev, mlx5, Tunnel Traffic
	<b>Discovered in Release:</b> 5.4-1.0.3.0
2566548	<b>Description:</b> On PPC systems when EEH is enabled, running fw sync reset (either by mlx5fwreset with flag --sync 1 or by devlink dev reload action fw_activate), the EEH may catch the PCI reset and take ownership on the flow. When run few times in sequence, the EEH may also decide to disable the device.
	<b>Workaround:</b> Administrator may disable EEH before running firmware sync reset on the device.
	<b>Keywords:</b> PPC, EEH
	<b>Discovered in Release:</b> 5.4-1.0.3.0
2617950	<b>Description:</b> TX port timestamp feature is supported for kernel versions 3.15 and greater. On older kernel versions, the feature will not be supported and ptp_tx_* counters will not increment.

Internal Ref. Number	Issue
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Ethtool
	<b>Discovered in Release:</b> 5.4-1.0.3.0
2390731	<b>Description:</b> Ethtool does not display Port Speed advertised/capability above 100Gb/s over and below kernels 5.0, even when supported.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Ethtool, Port Speed
	<b>Discovered in Release:</b> 5.4-1.0.3.0

Internal Ref. Number	Issue
2687198	<b>Description:</b> Activating VF/SF LAG when at least one VF/SF is still bound may lead to an internal error in the firmware.
	<b>Workaround:</b> Make sure all VFs/SFs are unbound prior to VF/SF LAG activation/deactivation.
	<b>Keywords:</b> VF, SF, Firmware, Binding
	<b>Discovered in Release:</b> 5.4-1.0.3.0

Internal Ref. Number	Issue
2585575	<b>Description:</b> After disabling sync reset by setting enable_remote_dev_reset to false, running firmware sync reset a few times may lead to general protection fault and system may get stuck.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Firmware Upgrade
	<b>Discovered in Release:</b> 5.3-1.0.0.1
2582565	<b>Description:</b> Conducting a firmware reset or unbinding the PF while in switchdev mode may cause a kernel crash.
	<b>Workaround:</b> N/A

Internal Ref. Number	Issue
	<p><b>Keywords:</b> SwitchDev, ASAP<sup>2</sup>, Unbind, Firmware Reset</p> <p><b>Discovered in Release:</b> 5.3-1.0.0.1</p>
2587802	<p><b>Description:</b> PTP synchronization may be lost while using tx_port_ts private flag.</p> <p><b>Workaround:</b> Toggle private flag:  ethtool --set-priv-flags tx_port_ts off  ethtool --set-priv-flags tx_port_ts on  restart ptp4l application</p> <p><b>Keywords:</b> PTP Synchronization</p> <p><b>Discovered in Release:</b> 5.3-1.0.0.1</p>
2574943	<p><b>Description:</b> When running kernel 5.8 and bellow or RHEL 8.2 and below, sampled packets do not support tunnel information.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> ASAP<sup>2</sup>, sFLOW</p> <p><b>Discovered in Release:</b> 5.3-1.0.0.1</p>
2568417	<p><b>Description:</b> Upon upgrade to version 5.3, the package manager tool will install the new packages and then remove the old packages, a depmod WARNING on "mlx5_fpga_tools" will appear. This warning can be safely ignored. mlx5_fpga_tools is a module that existed in version 5.2 and was removed in 5.3.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Upgrade; mlx5_fpga_tools</p> <p><b>Discovered in Release:</b> 5.3-1.0.0.1</p>
2506425	<p><b>Description:</b> When installing kmod packages on EulerOS 2.0SP9 or OpenEuler 20.03, the following error appears: "modprobe: FATAL: could not get modversions of ". This error can be safely ignored. It is caused by incorrectly adding directories to a list of modules processed by /usr/sbin/weak-modules.</p> <p><b>Workaround:</b> N/A</p>

Internal Ref. Number	Issue
	<p><b>Keywords:</b> Installation; modules; kmod</p> <p><b>Discovered in Release:</b> 5.3-1.0.0.1</p>
2492509	<p><b>Description:</b> When installing the driver on OpenEuler or on EulerOS 2.0SP9, rebuilding the drivers (--add-kernel-support) with the --kmp option (to create kmod packages) generates packages that are uninstallable because they have a dependency on "/sbin/depmod" that the system does not provide. This dependency is created by a buggy kmod package building tool included with the distribution.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> add-kernel-support</p> <p><b>Discovered in Release:</b> 5.3-1.0.0.1</p>
2479327	<p><b>Description:</b> On SLES 12 SP5, if the kernel was upgraded to 4.12.14-122.46, it is not possible to rebuild kernel modules (--add-kernel-support) without upgrading gcc as well to at least 4.8.5-31.23.2.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Upgrade; SLES 12; add-kernel-support</p> <p><b>Discovered in Release:</b> 5.3-1.0.0.1</p>
2584441	<p><b>Description:</b> On SLES 12 SP5, if the kernel was upgraded to 4.12.14-122.46, it is not possible to rebuild kernel modules (--add-kernel-support) without upgrading gcc as well to at least 4.8.5-31.23.2.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Upgrade; SLES 12; add-kernel-support</p> <p><b>Discovered in Release:</b> 5.3-1.0.0.1</p>
2460865	<p><b>Description:</b> When setting MTU to low values, such as 68 bytes, packets may fail on oversize.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> MTU</p> <p><b>Discovered in Release:</b> 5.3-1.0.0.1</p>



Internal Ref. Number	Issue
2383318	<b>Description:</b> On kernels based on RedHat 7.2, the "tx_port_ts" feature, as set by ethtool <code>—set-priv-flags</code> , is disabled.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> RedHat; tx_port_ts
	<b>Discovered in Release:</b> 5.3-1.0.0.1
2575647	<b>Description:</b> An OvS-DPDK crash might occur while doing live-migration for VMs that use virtio-interfaces that are accelerated using OvS-DPDK vDPA ports.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> OvS-DPDK vDPA, Live-migration
	<b>Discovered in Release:</b> 5.3-1.0.0.1

Internal Ref. Number	Issue
2430071	<b>Description:</b> After reloading devlink in IPoIB setup, the IB link may stay in initialization state and require to run OpenSM to get the IB link to active state.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> IPoIB devlink reload
	<b>Discovered in Release:</b> 5.2-2.2.0.0
2302786	<b>Description:</b> On EulerOS 2.0 SP9 systems, the kernel ABI (kABI) between the base vhulk2006 kernel and the errata vhulk2008 kernel has been changed. It is now not possible to install MLNX_OFED compiled with KMP on vhulk2006 kernel on a vhulk2008 system.
	<b>Workaround:</b> Install MLNX_OFED with <code>--add-kernel-support</code> .
	<b>Keywords:</b> EulerOS; kABI; installation; <code>--add-kernel-support</code>
	<b>Discovered in Release:</b> 5.2-1.0.4.0
2398281	<b>Description:</b> A crash in the TLS Rx socket cleanup flow may occur due to a kernel issue where a wrong extra call to <code>tls_dev_del</code> is made.

Internal Ref. Number	Issue
	<p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> TLS RX device offload</p> <p><b>Discovered in Release:</b> 5.2-1.0.4.0</p>
2407415	<p><b>Description:</b> OpenEuler 20.03 Aarch64 with errata kernels 4.19.90-2011.6.0.0049.oe1.aarch64 and 4.19.90-2012.5.0.0054.oe1.aarch64 are incompatible with MLNX_OFED kmod-mlnx-ofa_kernel.</p> <p><b>Workaround:</b> Install MLNX_OFED with --add-kernel-support.</p> <p><b>Keywords:</b> OpenEuler; Aarch64; installation; --add-kernel-support</p> <p><b>Discovered in Release:</b> 5.2-1.0.4.0</p>
2348077	<p><b>Description:</b> RDMA device name for VFs may change after resetting all VFs at once.</p> <p><b>Workaround:</b> Either reset interfaces one by one with a delay in between, or use a network interface naming scheme with predictable interface names, such as NAME_PCI or NAME_GUID. Copy /lib/udev/rules.d/60-rdma-persistent-naming.rules to /etc/udev/rules.d/ and edit the last line accordingly. Note that this will change interface names.</p> <p><b>Keywords:</b> RDMA; VF</p> <p><b>Discovered in Release:</b> 5.2-1.0.4.0</p>
2381713	<p><b>Description:</b> esp4_offload and esp6_offload modules are expected to be loaded according to the list determined by the default kernel. However, these modules cannot be loaded when working over Debian 10 with non-default custom kernel as they are not included in it.</p> <p><b>Workaround:</b> Either install MLNX_OFED using --add-kernel-support, or rebuild the non-default custom kernel to include these modules.</p> <p><b>Keywords:</b> esp4_offload; esp6_offload; kernel, Debian</p> <p><b>Discovered in Release:</b> 5.2-1.0.4.0</p>
2382898	<p><b>Description:</b> On kernel 4.14, there is no traffic for UDP or TCP with payload size larger than 1398 on GENEVE IPv6 over VLAN tag interface.</p> <p><b>Workaround:</b> N/A</p>

Internal Ref. Number	Issue
	<p><b>Keywords:</b> GENEVE; stag; VLAN; UDP</p> <p><b>Discovered in Release:</b> 5.2-1.0.4.0</p>
2326155	<p><b>Description:</b> When toggling the link state while running RoCE traffic, the below warning may appear in the dmesg: __ib_cache_gid_add: unable to add gid &lt;gid&gt; error=-28</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> RoCE; __ib_cache_gid_add</p> <p><b>Discovered in Release:</b> 5.2-1.0.4.0</p>
2329654	<p><b>Description:</b> Running XDP over an IP tunnel may fail when working with kernels as old as version 4.14.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> XDP, Kernel</p> <p><b>Discovered in Release:</b> 5.2-1.0.4.0</p>
2249156	<p><b>Description:</b> MLNX_OFED installation will remove qperf package in case it was done after qperf installation.</p> <p><b>Workaround:</b> Make sure to install qperf package after installing MLNX_OFED, or re-install qperf after installing MLNX_OFED.</p> <p><b>Keywords:</b> Installation; qperf</p> <p><b>Discovered in Release:</b> 5.2-1.0.4.0</p>
2355956	<p><b>Description:</b> OFED installation requires kernel config CONFIG_DEBUG_INFO to be set.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Installation; CONFIG_DEBUG_INFO</p> <p><b>Discovered in Release:</b> 5.2-1.0.4.0</p>
2362781	<p><b>Description:</b> Openibd may fail to unload the Inbox driver mlx5_ib on Ubuntu 18.04 PPC Boston server due to a bug in the Inbox drivers.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Openibd; Inbox; Ubuntu; mlx5_ib</p>

Internal Ref. Number	Issue
	<b>Discovered in Release:</b> 5.2-1.0.4.0
2367659	<p><b>Description:</b> Upgrading the MLNX_OFED version that is configured as a YUM repository may yield warning messages from depmod about unknown symbols, such as:  depmod: WARNING: /lib/modules/4.18.0-240.el8.x8664/extra/iser/ib_iser.ko needs unknown symbol ib_fmr_pool_unmap  depmod: WARNING: /lib/modules/4.18.0-240.el8.x8664/extra/srp/ib_srp.ko needs unknown symbol ib_create_qp_user  These warnings appear since the RPM packages upgrade occurs sequentially, and there is an upgrade dependency between some of the modules, which would create a state of upgrade inconsistency. These warnings are temporary and can be ignored as eventually all modules will be upgraded, and the warnings will no longer appear.</p>
	<b>Workaround:</b> N/A
	<b>Keywords:</b> YUM; RPM; symbol; depmod; ISER; SRP
	<b>Discovered in Release:</b> 5.2-1.0.4.0
2385269	<p><b>Description:</b> The number of connections offloaded is limited to 100K when working with Kernel v5.9.</p>
	<b>Workaround:</b> N/A
	<b>Keywords:</b> ASAP2; Connection Tracking; Kernel
	<b>Discovered in Release:</b> 5.2-1.0.4.0
2393169	<p><b>Description:</b> Mirroring is not supported with Connection Tracking when the source port is a VxLAN device.</p>
	<b>Workaround:</b> N/A
	<b>Keywords:</b> ASAP <sup>2</sup> ; Connection Tracking; Mirroring
	<b>Discovered in Release:</b> 5.2-1.0.4.0
2395082	<p><b>Description:</b> A call trace may take place when moving from SwitchDev mode back to Legacy mode in Kernel v5.9 due to a kernel issue in tcf_block_unbind.</p>
	<b>Workaround:</b> N/A
	<b>Keywords:</b> ASAP <sup>2</sup> ; SwitchDev; call trace; kernel; tcf_block_unbind

Internal Ref. Number	Issue
	<b>Discovered in Release:</b> 5.2-1.0.4.0

Internal Ref. Number	Issue
2354899	<b>Description:</b> ODP is not supported on RHEL7.x systems when running over an ETH link layer with RoCE disabled.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> ODP, RHEL, RoCE
	<b>Discovered in Release:</b> 5.1-2.5.8.0
2338150	<b>Description:</b> Scatter to <b>CQE</b> feature should be disabled for the GPUDirect tests to work.
	<b>Workaround:</b> Set the MLX5_SCATTER_TO_CQE environment variable to 0 before the <code>ib_send_bw</code> command. For example: MLX5_SCATTER_TO_CQE=0 <code>ib_send_bw -d &lt;...&gt;</code>
	<b>Keywords:</b> CQE, GPUDirect
	<b>Discovered in Release:</b> 5.1-2.5.8.0
2295732	<b>Description:</b> Upgrading from legacy (mlnx-libs) to the current rdma-core based build using YUM (package manager) fails.
	<b>Workaround:</b> To perform this upgrade, either use the installer script or uninstall the old packages and install the new packages.
	<b>Keywords:</b> Legacy, mlnx-libs, rdma-core, installation
	<b>Discovered in Release:</b> 5.1-2.5.8.0
2295735	<b>Description:</b> Upgrading from legacy (mlnx-libs) to the current rdma-core based build using the apt-get (package manager) fails.
	<b>Workaround:</b> To perform this upgrade, either use the installer script or uninstall the old packages and install the new packages.
	<b>Keywords:</b> Legacy, mlnx-libs, rdma-core, apt, apt-get, installation
	<b>Discovered in Release:</b> 5.1-2.5.8.0

Internal Ref. Number	Issue
2248996	<b>Description:</b> Downgrading the firmware version for ConnectX-6 cards using "mlnx_ofed_install --fw-update-only --force-fw-update" fails.
	<b>Workaround:</b> Manually downgrade the firmware version - please see <a href="#">Firmware Update Instructions</a> .
	<b>Keywords:</b> Firmware, ConnectX-6
	<b>Discovered in Release:</b> 5.1-0.6.6.0
2175930	<b>Description:</b> When using OFED 5.1 on PPC architectures with kernels v5.5 or v5.6 and an old ethtool utility, a harmless warning call trace may appear in the dmesg due to mismatch between user space and kernel. The warning call trace mentions ethtool_notify.
	<b>Workaround:</b> Update the ethtool utility to version 5.6 on such systems in order to avoid the call trace.
	<b>Keywords:</b> PPC, ethtool_notify, kernel
	<b>Discovered in Release:</b> 5.1-0.6.6.0
2198764	<b>Description:</b> If MLNX_OFED is installed on a Debian or Ubuntu system that is run in chroot environment, the openibd service will not be enabled. If the chroot files are being used as a base of a full system, the openibd service is left disabled.
	<b>Workaround:</b> Currently, openibd is a sysv-init script that you can enable manually by running: update-rc.d openibd defaults
	<b>Keywords:</b> chroot, Debian , Ubuntu, openibd
	<b>Discovered in Release:</b> 5.1-0.6.6.0
2237134	<b>Description:</b> Running connection tracking (CT) with FW steering may cause CREATE_FLOW_TABLE command to fail with syndrome.
	<b>Workaround:</b> Configure OVS to use a single handler-thread: #ovs-vsctl set Open_vSwitch . other_config:n-handler-threads=1
	<b>Keywords:</b> Connection tracking, ASAP, OVS, FW steering
	<b>Discovered in Release:</b> 5.1-0.6.6.0
2239894	<b>Description:</b> Running OpenVSwitch offload with high traffic throughput can cause low insertion rate due to high CPU usage.

Internal Ref. Number	Issue
	<p><b>Workaround:</b> Reduce the number of combined channels of the uplink using "ethtool -L".</p> <p><b>Keywords:</b> Insertion rate, ASAP2</p> <p><b>Discovered in Release:</b> 5.1-0.6.6.0</p>
2240671	<p><b>Description:</b> Header rewrite action is not supported over RHEL/CentOS 7.4.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> ASAP, header rewrite, RHEL, RedHat, CentOS, OS</p> <p><b>Discovered in Release:</b> 5.1-0.6.6.0</p>
2242546	<p><b>Description:</b> Tunnel offload (encap/decap) may cause kernel panic if nf_tables module is not probed.</p> <p><b>Workaround:</b> Make sure to probe the nf_tables module before inserting any rule.</p> <p><b>Keywords:</b> Kernel v5.7, ASAP, kernel panic</p> <p><b>Discovered in Release:</b> 5.1-0.6.6.0</p>
2143007	<p><b>Description:</b> IPsec packets are dropped during heavy traffic due to a bug in net/xfrm Linux Kernel.</p> <p><b>Workaround:</b> Make sure the Kernel is modified to apply the following patch: "xfrm: Fix double ESP trailer insertion in IPsec crypto offload".</p> <p><b>Keywords:</b> IPsec, xfrm</p> <p><b>Discovered in Release:</b> 5.1-0.6.6.0</p>
2225952	<p><b>Description:</b> VF mirroring with TC policy skip_sw is not supported on RHEL/CentOS 7.4, 7.5 and 7.6 OSs.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> ASAP<sup>2</sup>, Mirroring, RHEL, RedHat, OS</p> <p><b>Discovered in Release:</b> 5.1-0.6.6.0</p>
2216521	<p><b>Description:</b> After upgrading MLNX_OFED from v5.0 or earlier, ibdev2netdev utility changes the installation prefix to /usr/sbin.</p>

Internal Ref. Number	Issue
	<p>Therefore, it cannot be found while found in the same SHELL environment.</p> <p><b>Workaround:</b> After installing MLNX_OFED, log out and log in again to refresh the SHELL environment.</p> <p><b>Keywords:</b> ibdev2netdev</p> <p><b>Discovered in Release:</b> 5.1-0.6.6.0</p>
2202520	<p><b>Description:</b> Rules with VLAN push/pop, encap/decap and header rewrite actions together are not supported.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> ASAP<sup>2</sup>, SwitchDev, VLAN push/pop, encap/decap, header rewrite</p> <p><b>Discovered in Release:</b> 5.1-0.6.6.0</p>
2210752	<p><b>Description:</b> Switching from Legacy mode to SwitchDev mode and vice-versa while TC rules exist on the NIC will result in failure.</p> <p><b>Workaround:</b> Before attempting to switch mode, make sure to delete all TC rules on the NIC or stop OpenvSwitch.</p> <p><b>Keywords:</b> ASAP<sup>2</sup>, Devlink, Legacy SR-IOV</p> <p><b>Discovered in Release:</b> 5.1-0.6.6.0</p>

Internal Ref. Number	Issue
2125036/2125031	<p><b>Description:</b> Upgrading the MLNX_OFED from an UPSTREAM_LIBS based version to an MLNX_LIBS based version fails unless the driver is uninstalled and then re-installed.</p> <p><b>Workaround:</b> Make sure to uninstall and re-install MLNX_OFED to complete the upgrade.</p>



Internal Ref. Number	Issue
	<p><b>Keywords:</b> Installation, UPSTREAM_LIBS, MLNX_LIBS</p> <p><b>Discovered in Release:</b> 5.0-2.1.8.0</p>
2105447	<p><b>Description:</b> hns_roce warning messages will appear in the dmesg after reboot on Euler2 SP3 OSs.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> hns_roce, dmesg, Euler</p> <p><b>Discovered in Release:</b> 5.0-2.1.8.0</p>
2110321	<p><b>Description:</b> Multiple driver restarts may cause IPoIB soft lockup.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Driver restart, IPoIB</p> <p><b>Discovered in Release:</b> 5.0-2.1.8.0</p>
2112251	<p><b>Description:</b> On kernels 4.10-4.14, when Geneve tunnel's remote endpoint is defined using IPv6, packets larger than MTU are not fragmented, resulting in no traffic sent.</p> <p><b>Workaround:</b> Define geneve tunnel's remote endpoint using IPv4.</p> <p><b>Keywords:</b> Kernel, Geneve, IPv4, IPv6, MTU, fragmentation</p> <p><b>Discovered in Release:</b> 5.0-2.1.8.0</p>
2102902	<p><b>Description:</b> A kernel panic may occur over RH8.0-4.18.0-80.el8.x86_64 OS when opening kTLS offload connection due to a bug in kernel TLS stack.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> TLS offload, mlx5e</p> <p><b>Discovered in Release:</b> 5.0-2.1.8.0</p>
2111534	<p><b>Description:</b> A Kernel panic may occur over Ubuntu19.04-5.0.0-38-generic OS when opening kTLS offload connection due to a bug in the Kernel TLS stack.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> TLS offload, mlx5e</p>

Internal Ref. Number	Issue
	<b>Discovered in Release:</b> 5.0-2.1.8.0
2035950	<p><b>Description:</b> An internal error might take place in the firmware when performing any of the following in VF LAG mode, when at least one VF of either PF is still bound/attached to a VM.</p> <ol style="list-style-type: none"> <li>1. Removing PF from the bond (using ifdown, ip link or any other function)</li> <li>2. Attempting to disable SR-IOV</li> </ol>
	<b>Workaround:</b> N/A
	<b>Keywords:</b> VF LAG, binding, firmware, FW, PF, SR-IOV
	<b>Discovered in Release:</b> 5.0-1.0.0.0
2044544	<p><b>Description:</b> When working with OSs with Kernel v4.10, bonding module does not allow setting MTUs larger than 1500 on a bonding interface.</p>
	<b>Workaround:</b> Upgrade your Kernel version to v4.11 or above.
	<b>Keywords:</b> Bonding, MTU, Kernel
	<b>Discovered in Release:</b> 5.0-1.0.0.0
1882932	<p><b>Description:</b> Libibverbs dependencies are removed during OFED installation, requiring manual installation of libraries that OFED does not reinstall.</p>
	<b>Workaround:</b> Manually install missing packages.
	<b>Keywords:</b> libibverbs, installation
	<b>Discovered in Release:</b> 5.0-1.0.0.0
2058535	<p><b>Description:</b> ibdev2netdev command returns duplicate devices with different ports in SwitchDev mode.</p>
	<b>Workaround:</b> Use /opt/mellanox/iproute2/sbin/rdma link show command instead.
	<b>Keywords:</b> ibdev2netdev
	<b>Discovered in Release:</b> 5.0-1.0.0.0

Internal Ref. Number	Issue
2072568	<b>Description:</b> In RHEL/CentOS 7.2 OSs, adding drop rules when act_gact is not loaded may cause a kernel crash.
	<b>Workaround:</b> Preload all needed modules to avoid such a scenario (cls_flower, act_mirred, act_gact, act_tunnel_key and act_vlan).
	<b>Keywords:</b> RHEL/CentOS 7.2, Kernel 4.9, call trace, ASAP
	<b>Discovered in Release:</b> 5.0-1.0.0.0
2093698	<b>Description:</b> VF LAG configuration is not supported when the NUM_OF_VFS configured in mlxconfig is higher than 64.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> VF LAG, SwitchDev mode, ASAP
	<b>Discovered in Release:</b> 5.0-1.0.0.0
2093746	<b>Description:</b> Devlink health dumps are not supported on kernels lower than v5.3.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Devlink, health report, dump
	<b>Discovered in Release:</b> 5.0-1.0.0.0
2000590	<b>Description:</b> Sending packets larger than MTU is not supported when working with OVS-DPDK.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> MTU, OVS-DPDK
	<b>Discovered in Release:</b> 5.0-1.0.0.0
2062900	<b>Description:</b> Moving VF from SwitchDev mode to Legacy mode while the representor is being used by OVS-DPDK results in a segmentation fault.
	<b>Workaround:</b> To move VF to Legacy mode with no error, make sure to delete the ports from the OVS.
	<b>Keywords:</b> SwitchDev, Legacy, representor, OVS-DPDK
	<b>Discovered in Release:</b> 5.0-1.0.0.0

Internal Ref. Number	Issue
2075942	<b>Description:</b> Huge pages configuration is lost each time the server is configured.
	<b>Workaround:</b> Re-configure the huge pages after each reboot, or configure them as a kernel parameter.
	<b>Keywords:</b> Huge pages, reboot, OVS-DPDK
	<b>Discovered in Release:</b> 5.0-1.0.0.0
2067012	<b>Description:</b> MLNX_OFED cannot be installed on Debian 9.11 OS in SwitchDev mode.
	<b>Workaround:</b> Install OFED with the flag --add-kernel-support.
	<b>Keywords:</b> ASAP, SwitchDev, Debian, Kernel
	<b>Discovered in Release:</b> 5.0-1.0.0.0
2036572	<b>Description:</b> When using a thread domain and the lockless rdma-core ibv_post_send path, there is an additional CPU penalty due to required barriers around the device MMIO buffer that were omitted in MLNX_OFED.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> rdma-core, write-combining, MMIO buffer
	<b>Discovered in Release:</b> 5.0-1.0.0.0

Internal Ref. Number	Issue
-	<b>Description:</b> The argparse module is installed by default in Python versions =>2.7 and >=3.2. In case an older Python version is used, the argparse module is not installed by default.
	<b>Workaround:</b> Install the argparse module manually.
	<b>Keywords:</b> Python, MFT, argparse, installation

Internal Ref. Number	Issue
	<b>Discovered in Release:</b> 4.7-3.2.9.0
1997230	<b>Description:</b> Running mlxfwreset or unloading mlx5_core module while kontrak flows are offloaded may cause a call trace in the kernel.
	<b>Workaround:</b> Stop OVS service before calling mlxfwreset or unloading mlx5_core module.
	<b>Keywords:</b> Kontrak, ASAP, OVS, mlxfwrest, unload
	<b>Discovered in Release:</b> 4.7-3.2.9.0
1955352	<b>Description:</b> Moving 2 ports to SwitchDev mode in parallel is not supported.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> ASAP, SwitchDev
	<b>Discovered in Release:</b> 4.7-3.2.9.0
1979958	<b>Description:</b> VxLAN IPv6 offload is not supported over CentOS/RHEL v7.2 OSs.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Tunnel, VXLAN, ASAP, IPv6
	<b>Discovered in Release:</b> 4.7-3.2.9.0
1991710	<b>Description:</b> PRIO_TAG_REQUIRED_EN configuration is not supported and may cause call trace.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> ASAP, PRIO_TAG, mstconfig
	<b>Discovered in Release:</b> 4.7-3.2.9.0
1967866	<b>Description:</b> Enabling ECMP offload requires the VFs to be unbound and VMs to be shut down.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> ECMP, Multipath, ASAP <sup>2</sup>
	<b>Discovered in Release:</b> 4.7-3.2.9.0

Internal Ref. Number	Issue
1921981	<b>Description:</b> On Ubuntu, Debian and RedHat 8 and above OSS, parsing the mfa2 file using the mstarchive might result in a segmentation fault.
	<b>Workaround:</b> Use mlxarchive to parse the mfa2 file instead.
	<b>Keywords:</b> MFT, mfa2, mstarchive, mlxarchive, Ubuntu, Debian, RedHat, operating system
	<b>Discovered in Release:</b> 4.7-1.0.0.1
1840288	<b>Description:</b> MLNX_OFED does not support XDP features on RedHat 7 OS, despite the declared support by RedHat.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> XDP, RedHat
	<b>Discovered in Release:</b> 4.7-1.0.0.1
1821235	<b>Description:</b> When using mlx5dv_dr API for flow creation, for flows which execute the "encapsulation" action or "push vlan" action, metadata C registers will be reset to zero.
	<b>Workaround:</b> Use the both actions at the end of the flow process.
	<b>Keywords:</b> Flow steering
	<b>Discovered in Release:</b> 4.7-1.0.0.1

Internal Ref. Number	Issue
1504785	<b>Description:</b> A lost interrupt issue in pass-through virtual machines may prevent the driver from loading, followed by printing managed pages errors to the dmesg.
	<b>Workaround:</b> Restart the driver.
	<b>Keywords:</b> VM, virtual machine
	<b>Discovered in Release:</b> 4.6-1.0.1.1
1764415	<b>Description:</b> Unbinding PFs on LAG devices results in a "Failed to modify QP to RESET" error message.
	<b>Workaround:</b> N/A

Internal Ref. Number	Issue
	<p><b>Keywords:</b> RoCE LAG, unbind, PF, RDMA</p> <p><b>Discovered in Release:</b> 4.6-1.0.1.1</p>
1806565	<p><b>Description:</b> RoCE default GIDs v1 and v2 are derived from the MAC address of the corresponding netdevice's PCI function, and they resemble the IPv6 address. However, in systems where the IPv6 link local address generated does not depend on the MAC address, RoCEv2 default GID should not be used.</p> <p><b>Workaround:</b> Use RoCEv2 default GID.</p> <p><b>Keywords:</b> RoCE</p> <p><b>Discovered in Release:</b> 4.6-1.0.1.1</p>
-	<p><b>Description:</b> Aging is not functional on bond device in RHEL 7.6.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> VF LAG, ASAP<sup>2</sup></p> <p><b>Discovered in Release:</b> 4.6-1.0.1.1</p>
1747774	<p><b>Description:</b> In VF LAG mode, outgoing traffic in load balanced mode is according to the origin ring, thus, half of the rings will be coupled with port 1 and half with port 2. All the traffic on the same ring will be sent from the same port.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> VF LAG, ASAP<sup>2</sup></p> <p><b>Discovered in Release:</b> 4.6-1.0.1.1</p>
1753629	<p><b>Description:</b> A bonding bug found in Kernels 4.12 and 4.13 may cause a slave to become permanently stuck in BOND_LINK_FAIL state. As a result, the following message may appear in dmesg: bond: link status down for interface eth1, disabling it in 100 ms</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Bonding, slave</p> <p><b>Discovered in Release:</b> 4.6-1.0.1.1</p>

Internal Ref. Number	Issue
1712068	<b>Description:</b> Uninstalling MLNX_OFED automatically results in the uninstallation of several libraries that are included in the MLNX_OFED package, such as InfiniBand-related libraries.
	<b>Workaround:</b> If these libraries are required, reinstall them using the local package manager (yum/dnf).
	<b>Keywords:</b> MLNX_OFED libraries
	<b>Discovered in Release:</b> 4.6-1.0.1.1
-	<b>Description:</b> Due to changes in libraries, MFT v4.11.0 and below are not forward compatible with MLNX_OFED v4.6-1.0.0.0 and above. Therefore, with MLNX_OFED v4.6-1.0.0.0 and above, it is recommended to use MFT v4.12.0 and above.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> MFT compatible
	<b>Discovered in Release:</b> 4.6-1.0.1.1
1730840	<b>Description:</b> On ConnectX-4 HCAs, GID index for RoCE v2 is inconsistent when toggling between enabled and disabled interface modes.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> RoCE v2, GID
	<b>Discovered in Release:</b> 4.6-1.0.1.1
1717428	<b>Description:</b> On kernels 4.10-4.14, MTUs larger than 1500 cannot be set for a GRE interface with any driver (IPv4 or IPv6).
	<b>Workaround:</b> Upgrade your kernel to any version higher than v4.14.
	<b>Keywords:</b> Fedora 27, gretap, ip_gre, ip_tunnel, ip6_gre, ip6_tunnel
	<b>Discovered in Release:</b> 4.6-1.0.1.1
1748343	<b>Description:</b> Driver reload takes several minutes when a large number of VFs exists.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> VF, SR-IOV



Internal Ref. Number	Issue
	<b>Discovered in Release:</b> 4.6-1.0.1.1
1733974	<b>Description:</b> Running heavy traffic (such as 'ping flood') while bringing up and down other mlx5 interfaces may result in "INFO: rcu_preempt dectected stalls on CPUS/tasks:" call traces.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> mlx5
	<b>Discovered in Release:</b> 4.6-1.0.1.1
-	<b>Description:</b> On ConnectX-6 HCAs and above, an attempt to configure advertisement (any bitmap) will result in advertising the whole capabilities.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> 200Gb/s, advertisement, Ethtool
	<b>Discovered in Release:</b> 4.6-1.0.1.1

Internal Ref. Number	Issue
1699289	<b>Description:</b> HW LRO feature is disabled OOB, which results in increased CPU utilization on the Receive side. On ConnectX-5 adapter cards and above, this causes a bandwidth drop for a few streams.
	<b>Workaround:</b> Make sure to enable HW LRO in the driver: ethtool -k <intf> lro ethtool --set-priv-flag <intf> hw_lro on
	<b>Keywords:</b> HW LRO, ConnectX-5 and above
	<b>Discovered in Release:</b> 4.5-1.0.1.0
1403313	<b>Description:</b> Attempting to allocate an excessive number of VFs per PF in operating systems with kernel versions below v4.15 might fail due to a known issue in the Kernel.

Internal Ref. Number	Issue
	<p><b>Workaround:</b> Make sure to update the Kernel version to v4.15 or above.</p> <p><b>Keywords:</b> VF, PF, IOMMU, Kernel, OS</p> <p><b>Discovered in Release:</b> 4.5-1.0.1.0</p>
-	<p><b>Description:</b> NEO-Host is not supported on the following OSs:</p> <ul style="list-style-type: none"> <li>• SLES12 SP3</li> <li>• SLES12 SP4</li> <li>• SLES15</li> <li>• Fedora 28</li> <li>• RHEL7.1</li> <li>• RHEL7.4 ALT (Pegas1.0)</li> <li>• REL 7.5</li> <li>• RHEL7.6</li> <li>• XenServer 4.9</li> </ul> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> NEO-Host, operating systems</p> <p><b>Discovered in Release:</b> 4.5-1.0.1.0</p>
1521877	<p><b>Description:</b> On SLES 12 SP1 OSs, a kernel tracepoint issue may cause undefined behavior when inserting a kernel module with a wrong parameter.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> mlx5 driver, SLES 12 SP1</p> <p><b>Discovered in Release:</b> 4.5-1.0.1.0</p>

---

# User Manual

- [Introduction](#)
- [Installation](#)
- [Features Overview and Configuration](#)
- [Programming](#)
- [InfiniBand Fabric Utilities](#)
- [Troubleshooting](#)
- [Common Abbreviations and Related Documents](#)

## Introduction

This manual is intended for system administrators responsible for the installation, configuration, management and maintenance of the software and hardware of VPI (InfiniBand, Ethernet) adapter cards. It is also intended for application developers.

MLNX\_OFED is a single Virtual Protocol Interconnect (VPI) software stack which operates across all NVIDIA network adapter solutions supporting the following uplinks to servers:

Uplink/Adapter Card	Driver Name	Uplink Speed
BlueField-2	mlx5	<ul style="list-style-type: none"><li>• InfiniBand: SDR, FDR, EDR, HDR</li><li>• Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE<sup>2</sup>, 100GbE<sup>2</sup></li></ul>
BlueField		<ul style="list-style-type: none"><li>• InfiniBand: SDR, QDR, FDR, FDR10, EDR</li><li>• Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 100GbE</li></ul>

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

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


All NVIDIA network adapter cards are compatible with OpenFabrics-based RDMA protocols and software and are supported by major operating system distributions.

NVIDIA OFED is certified with the following products:

- NVIDIA Messaging Accelerator (VMA™) software: Socket acceleration library that performs OS bypass for standard socket-based applications. Please note, VMA support is provided separately from NVIDIA OFED support. For further information, please refer to the VMA documentation ([docs.nvidia.com/networking/category/vma](https://docs.nvidia.com/networking/category/vma)).
- NVIDIA Unified Fabric Manager (UFM®) software: Powerful platform for managing demanding scale-out computing fabric environments, built on top of the OpenSM industry standard routing engine.
- Fabric Collective Accelerator (FCA)—FCA is a NVIDIA MPI-integrated software package that utilizes CORE-Direct technology for implementing the MPI collectives communications.

## Stack Architecture

The figure below shows a diagram of the NVIDIA OFED stack, and how upper layer protocols (ULPs) interface with the hardware and with the kernel and userspace. The application level also shows the versatility of markets that NVIDIA OFED applies to.



The following subsections briefly describe the various components of the NVIDIA OFED stack.

### mlx4 VPI Driver

**i Note**

This driver is no longer supported in MLNX\_OFED. To work with ConnectX-3® and ConnectX-3 Pro NICs, please refer to MLNX\_OFED LTS version available on the web.

### mlx5 Driver

mlx5 is the low-level driver implementation for the Connect-IB® and ConnectX-4 and above adapters designed by NVIDIA. ConnectX-4 and above adapter cards operate as a VPI adapter (Infiniband and Ethernet). The mlx5 driver is comprised of the following kernel modules:

## **Note**

Please note that Connect-IB card is no longer supported in MLNX\_OFED. To work with this card, please refer to MLNX\_OFED LTS version available on the web.

### **mlx5\_core**

Acts as a library of common functions (e.g. initializing the device after reset) required by ConnectX-4 and above adapter cards. mlx5\_core driver also implements the Ethernet interfaces for ConnectX-4 and above. mlx5 drivers do not require the mlx5\_en module as the Ethernet functionalities are built-in in the mlx5\_core module.

### **mlx5\_ib**

Handles InfiniBand-specific functions and plugs into the InfiniBand mid layer.

### **libmlx5**

libmlx5 is the provider library that implements hardware specific user-space functionality. If there is no compatibility between the firmware and the driver, the driver will not load and a message will be printed in the dmesg.

The following are the libmlx5 **Legacy** and **RDMA-Core** environment variables:

- MLX5\_FREEZE\_ON\_ERROR\_CQE
  - Causes the process to hang in a loop of completion with error, which is not flushed with error or retry exceeded occurs/
  - Otherwise disabled
- MLX5\_POST\_SEND\_PREFER\_BF
  - Configures every work request that can use blue flame will use blue flame
- Otherwise, blue flame depends on the size of the message and inline indication in the packet

- MLX5\_SHUT\_UP\_BF
  - Disables blue flame feature
  - Otherwise, do not disable
- MLX5\_SINGLE\_THREADED
  - All spinlocks are disabled
  - Otherwise, spinlocks enabled
  - Used by applications that are single threaded and would like to save the overhead of taking spinlocks.
- MLX5\_CQE\_SIZE
  - 64—completion queue entry size is 64 bytes (default)
  - 128—completion queue entry size is 128 bytes
- MLX5\_SCATTER\_TO\_CQE
  - Small buffers are scattered to the completion queue entry and manipulated by the driver. Valid for RC transport.
  - Default is 1, otherwise disabled

The following are libmlx5 **Legacy** only environment variables:

- MLX5\_ENABLE\_CQE\_COMPRESSION
  - Saves PCIe bandwidth by compressing a few CQEs into a smaller amount of bytes on PCIe. Setting this variable enables CQE compression.
  - Default value 0 (disabled)
- MLX5\_RELAXED\_PACKET\_ORDERING\_ON  
See “[Out-of-Order \(OOO\) Data Placement](#)” section.

## Mid-layer Core



Core services include management interface (MAD), connection manager (CM) interface, and Subnet Administrator (SA) interface. The stack includes components for both user-mode and kernel applications. The core services run in the kernel and expose an interface to user-mode for verbs, CM and management.

## **Upper Layer Protocols (ULPs)**

### **IP over IB (IPoIB)**

The IP over IB (IPoIB) driver is a network interface implementation over InfiniBand. IPoIB encapsulates IP datagrams over an InfiniBand connected or datagram transport service. IPoIB pre-appends the IP datagrams with an encapsulation header and sends the outcome over the InfiniBand transport service. The transport service is Unreliable Datagram (UD) by default, but it may also be configured to be Reliable Connected (RC), in case RC is supported. The interface supports unicast, multicast and broadcast. For details, see [“IP over InfiniBand \(IPoIB\)”](#) section.

### **iSCSI Extensions for RDMA (iSER)**

iSCSI Extensions for RDMA (iSER) extends the iSCSI protocol to RDMA. It permits data to be transferred directly into and out of SCSI buffers without intermediate data copies. For further information, please refer to [“iSCSI Extensions for RDMA \(iSER\)”](#) section.

### **SCSI RDMA Protocol (SRP)**

SCSI RDMA Protocol (SRP) is designed to take full advantage of the protocol offload and RDMA features provided by the InfiniBand architecture. SRP allows a large body of SCSI software to be readily used on InfiniBand architecture. The SRP driver—known as the SRP Initiator—differs from traditional low-level SCSI drivers in Linux. The SRP Initiator does not control a local HBA; instead, it controls a connection to an I/O controller—known as the SRP Target—to provide access to remote storage devices across an InfiniBand fabric. The SRP Target resides in an I/O unit and provides storage services. See [“SRP—SCSI RDMA Protocol”](#) section.

### **User Direct Access Programming Library (uDAPL)**

User Direct Access Programming Library (uDAPL) is a standard API that promotes data center application data messaging performance, scalability, and reliability over RDMA interconnects InfiniBand and RoCE. The uDAPL interface is defined by the DAT collaborative. This release of the uDAPL reference implementation package for both DAT 1.2 and 2.0 specification is timed to coincide with OFED release of the Open Fabrics ([openfabrics.org](http://openfabrics.org)) software stack.

## MPI

Message Passing Interface (MPI) is a library specification that enables the development of parallel software libraries to utilize parallel computers, clusters, and heterogeneous networks. NVIDIA OFED includes the following MPI implementation over InfiniBand:

- Open MPI – an open source MPI-2 implementation by the Open MPI Project

NVIDIA OFED also includes MPI benchmark tests such as OSU BW/LAT, Intel MPI BeBenchmark and Presta.

## InfiniBand Subnet Manager

All InfiniBand-compliant ULPs require a proper operation of a Subnet Manager (SM) running on the InfiniBand fabric, at all times. An SM can run on any node or on an IB switch. OpenSM is an InfiniBand-compliant Subnet Manager, and it is installed as part of NVIDIA OFED<sup>1</sup>.

1. OpenSM is disabled by default. See "[NVIDIA SM](#)" section for details on enabling it.

## Diagnostic Utilities

NVIDIA OFED includes the following two diagnostic packages for use by network and data center managers:

- ibutils—NVIDIA diagnostic utilities
- infiniband-diags—OpenFabrics Alliance InfiniBand diagnostic tools

## NVIDIA Firmware Tools (MFT)

The NVIDIA Firmware Tools package is a set of firmware management tools for a single InfiniBand node. MFT can be used for:

- Generating a standard or customized NVIDIA firmware image
- Burning a firmware image to a single InfiniBand node

MFT includes a set of tools used for performing firmware update and configuration, as well as debug and diagnostics, and provides MST service. For the full list of available tools within MFT, please refer to MFT documentation ([docs.nvidia.com/networking/category/mft](https://docs.nvidia.com/networking/category/mft)).

# Package Contents

## ISO Image

NVIDIA OFED for Linux (MLNX\_OFED\_LINUX) is provided as ISO images or as a tarball, one per supported Linux distribution and CPU architecture, that includes *source code* and *binary* RPMs, firmware, utilities, and documentation. The ISO image contains an installation script (called `mlnxofedinstall`) that performs the necessary steps to accomplish the following:

- Discover the currently installed kernel
- Uninstall any InfiniBand stacks that are part of the standard operating system distribution or another vendor's commercial stack
- Install the MLNX\_OFED\_LINUX binary RPMs (if they are available for the current kernel)
- Identify the currently installed InfiniBand HCAs and perform the required firmware updates

## Software Components

MLNX\_OFED\_LINUX contains the following software components:

- NVIDIA Host Channel Adapter Drivers
  - `mlx5`
    - `mlx5_ib`
    - `mlx5_core` (includes Ethernet)
- Mid-layer core
  - Verbs, MADs, SA, CM, CMA, uVerbs, uMADs
- Upper Layer Protocols (ULPs)

- IPoIB, SRP Initiator and SRP
- MPI
  - Open MPI stack supporting the InfiniBand, RoCE and Ethernet interfaces
  - MPI benchmark tests (OSU BW/LAT, Intel MPI Benchmark, Presta)
- OpenSM: InfiniBand Subnet Manager
- Utilities
  - Diagnostic tools
  - Performance tests
  - Sysinfo (see [Sysinfo User Manual](#))
- Firmware tools (MFT)
- Source code for all the OFED software modules (for use under the conditions mentioned in the modules' LICENSE files)
- Documentation

## Firmware

The ISO image includes the following firmware item:

- mlnx-fw-updater RPM/DEB package, which contains firmware binaries for supported devices (using mlxfwmanager tool).

## Directory Structure

The ISO image of MLNX\_OFED\_LINUX contains the following files and directories:

- mlnxfedinstall—the MLNX\_OFED\_LINUX installation script.

- `ofed_uninstall.sh`—This is the `MLNX_OFED_LINUX` un-installation script.
- `<RPMS folders>`—Directory of binary RPMs for a specific CPU architecture.
- `src/`—Directory of the OFED source tarball.

### **Note**

`MLNX_OFED` includes the OFED source RPM packages used as a build platform for kernel code but does not include the sources of NVIDIA proprietary packages.

- `mlnx_add_kernel_support.sh`—Script required to rebuild `MLNX_OFED_LINUX` for customized kernel version on supported Linux Distribution
- RPM based—A script required to rebuild `MLNX_OFED_LINUX` for customized kernel version on supported RPM-based Linux Distribution
- `docs/`—Directory of NVIDIA OFED related documentation

## Module Parameters

### `mlx5_core` Module Parameters

The `mlx5_core` module supports a single parameter used to select the profile which defines the number of resources supported.

<i>prof_sel</i>	<p>The parameter name for selecting the profile. The supported values for profiles are:</p> <ul style="list-style-type: none"> <li>• 0 - for medium resources, medium performance</li> <li>• 1 - for low resources</li> <li>• 2 - for high performance (int) (default)</li> </ul>
-----------------	---

	<ul style="list-style-type: none"> <li>• 3 - fast unload</li> <li>• 4 - without DCT enable</li> </ul>
guids	charp
node_guid	guids configuration. This module parameter will be obsolete!
debug_mask	debug_mask: 1 = dump cmd data, 2 = dump cmd exec time, 3 = both. Default=0 (uint)
probe_vf	probe VFs or not, 0 = not probe, 1 = probe. Default = 1 (bool)
num_of_groups	Controls the number of large groups in the FDB flow table. Default=4; Range=1-1024

### ib\_core Parameters

send_queue_size	Size of send queue in number of work requests (int)
recv_queue_size	Size of receive queue in number of work requests (int)
force_mr	Force usage of MRs for RDMA READ/WRITE operations (bool)
roce_v1_noncompat_gid	Default GID auto configuration (Default: yes) (bool)

### ib\_ipoib Parameters

max_nonsrq_conn_qp	Max number of connected-mode QPs per interface (applied only if shared receive queue is not available) (int)
mcast_debug_level	Enable multicast debug tracing if > 0 (int)
send_queue_size	Number of descriptors in send queue (int)
recv_queue_size	Number of descriptors in receive queue (int)
debug_level	Enable debug tracing if > 0 (int)
ipoib_enhanced	Enable IPoIB enhanced for capable devices (default = 1) (0-1) (int)

# Device Capabilities

Normally, an application needs to query the device capabilities before attempting to create a resource. It is essential for the application to be able to operate over different devices with different capabilities.

Specifically, when creating a QP, the user needs to specify the maximum number of outstanding work requests that the QP supports. This value should not exceed the queried capabilities. However, even when you specify a number that does not exceed the queried capability, the verbs can still fail since some other factors such as the number of scatter/gather entries requested, or the size of the inline data required, affect the maximum possible work requests. Hence an application should try to decrease this size (halving is a good new value) and retry until it succeeds.

# Installation

This chapter describes how to install and test the NVIDIA OFED for Linux package on a single host machine with NVIDIA InfiniBand and/or Ethernet adapter hardware installed.

The chapter contains the following sections:

- [Hardware and Software Requirements](#)
- [Downloading the Drivers](#)
- [Installing MLNX\\_OFED](#)
- [Uninstall](#)
- [Updating Firmware After Installation](#)
- [UEFI Secure Boot](#)
- [Performance Tuning](#)

# Features Overview and Configuration

## Note

It is recommended to enable the "above 4G decoding" BIOS setting for features that require a large amount of PCIe resources (e.g., SR-IOV with numerous VFs, PCIe Emulated Switch, Large BAR Requests).

The chapter contains the following sections:

- [Ethernet Network](#)
- [InfiniBand Network](#)
- [Storage Protocols](#)
- [Virtualization](#)
- [Resiliency](#)
- [Docker Containers](#)
- [HPC-X](#)
- [Fast Driver Unload](#)
- [OVS Offload Using ASAP<sup>2</sup> Direct](#)

## Programming



## **Note**

This chapter is aimed for application developers and expert users that wish to develop applications over MLNX\_OFED.

## **Raw Ethernet Programming**

Raw Ethernet programming enables writing an application that bypasses the kernel stack. To achieve this, packet headers and offload options need to be provided by the application.

For a basic example on how to use Raw Ethernet programming, refer to the [Raw Ethernet Programming: Basic Introduction—Code Example](#) Community post.

## **Packet Pacing**

Packet pacing is a raw Ethernet sender feature that enables controlling the rate of each QP, per send queue.

For a basic example on how to use packet pacing per flow over libibverbs, refer to [Raw Ethernet Programming: Packet Pacing—Code Example](#) Community post.

## **TCP Segmentation Offload (TSO)**

TCP Segmentation Offload (TSO) enables the adapter cards to accept a large amount of data with a size greater than the MTU size. The TSO engine splits the data into separate packets and inserts the user-specified L2/L3/L4 headers automatically per packet. With the usage of TSO, CPU is offloaded from dealing with a large throughput of data.

To be able to program that on the sender side, refer to the [Raw Ethernet Programming: TSO—Code Example](#) Community post.

## **ToS Based Steering**

ToS/DSCP is an 8-bit field in the IP packet that enables different service levels to be assigned to network traffic. This is achieved by marking each packet in the network with a DSCP code and appropriating the corresponding level of service to it.

To be able to steer packets according to the ToS field on the receiver side, refer to the [Raw Ethernet Programming: ToS—Code Example](#) Community post.

## Flow ID Based Steering

Flow ID based steering enables developing a code that will steer packets using flow ID when developing Raw Ethernet over verbs. For more information on flow ID based steering, refer to the [Raw Ethernet Programming: Flow ID Steering—Code Example](#) Community post.

## VXLAN Based Steering

VXLAN based steering enables developing a code that will steer packets using the VXLAN tunnel ID when developing Raw Ethernet over verbs. For more information on VXLAN based steering, refer to the [Raw Ethernet Programming: VXLAN Steering—Code Example](#) Community post.

## Device Memory Programming

### Note

This feature is supported on ConnectX-5/ConnectX-5 Ex adapter cards and above only.

Device Memory is an API that allows using on-chip memory located on the device as a data buffer for send/receive and RDMA operations. The device memory can be mapped and accessed directly by user and kernel applications, and can be allocated in various

sizes, registered as memory regions with local and remote access keys for performing the send/receive and RDMA operations.

Using the device memory to store packets for transmission can significantly reduce transmission latency compared to the host memory.

## Device Memory Programming Model

The new API introduces a similar procedure to the host memory for sending packets from the buffer:

- `ibv_alloc_dm()/ibv_free_dm()` - to allocate/free device memory
- `ibv_reg_dm_mr` - to register the allocated device memory buffer as a memory region and get a memory key for local/remote access by the device
- `ibv_memcpy_to_dm` - to copy data to a device memory buffer
- `ibv_memcpy_from_dm` - to copy data from a device memory buffer
- `ibv_post_send/ibv_post_receive` - to request the device to perform a send/receive operation using the memory key

For examples, see [Device Memory](#).

## RDMA-CM QP Timeout Control

RDMA-CM QP Timeout Control feature enables users to control the QP timeout for QPs created with RDMA-CM.

A new option in 'rdma\_set\_option' function has been added to enable overriding calculated QP timeout, in order to provide QP attributes for QP modification. To achieve that, `rdma_set_option()` should be called with the new flag `RDMA_OPTION_ID_ACK_TIMEOUT`.

Example:

```
rdma_set_option(cma_id, RDMA_OPTION_ID, RDMA_OPTION_ID_ACK_TIMEOUT,  
&timeout, sizeof(timeout));
```

## RDMA-CM Application Managed QP

Applications which do not create a QP through `rdma_create_qp()` may want to postpone the ESTABLISHED event on the passive side, to let the active side complete an application-specific connection establishment phase. For example, modifying the init state of the QP created by the application to RTR state, or make some preparations for receiving messages from the passive side. The feature returns a new event on the active side: `CONNECT_RESPONSE`, instead of `ESTABLISHED`, if `id->qp==NULL`. This gives the application a chance to perform the extra connection setup. Afterwards, the new `rdma_establish()` API should be called to complete the connection and generate an ESTABLISHED event on the passive side.

In addition, this feature exposes the `'rdma_init_qp_attr'` function in `librdmacm` API, which enables applications to get the parameters for creating Address Handler (AH) or control QP attributes after its creation.

## InfiniBand Fabric Utilities

This section first describes common configuration, interface, and addressing for all the tools in the package.

### Common Configuration, Interface and Addressing

#### Topology File (Optional)

An InfiniBand fabric is composed of switches and channel adapter (HCA/TCA) devices. To identify devices in a fabric (or even in one switch system), each device is given a GUID (a MAC equivalent). Since a GUID is a non-user-friendly string of characters, it is better to alias it to a meaningful, user-given name. For this objective, the IB Diagnostic Tools can be provided with a “topology file”, which is an optional configuration file specifying the IB fabric topology in user-given names.

For diagnostic tools to fully support the topology file, the user may need to provide the local system name (if the local hostname is not used in the topology file).

To specify a topology file to a diagnostic tool use one of the following two options:

1. On the command line, specify the file name using the option '-t <topology file name>'
2. Define the environment variable IBDIAG\_TOPO\_FILE

To specify the local system name to an diagnostic tool use one of the following two options:

1. On the command line, specify the system name using the option '-s <local system name>'
2. Define the environment variable IBDIAG\_SYS\_NAME

## InfiniBand Interface Definition

The diagnostic tools installed on a machine connect to the IB fabric by means of an HCA port through which they send MADs. To specify this port to an IB diagnostic tool use one of the following options:

1. On the command line, specify the port number using the option '-p <local port number>' (see below)
2. Define the environment variable IBDIAG\_PORT\_NUM

In case more than one HCA device is installed on the local machine, it is necessary to specify the device's index to the tool as well. For this use one of the following options:

1. On the command line, specify the index of the local device using the following option: '-i <index of local device>'
2. Define the environment variable IBDIAG\_DEV\_IDX

## Addressing

### Note

This section applies to the `ibdiagpath` tool only. A tool command may require defining the destination device or port to which it applies.

The following addressing modes can be used to define the IB ports:

- Using a Directed Route to the destination: (Tool option `'-d'`)  
This option defines a directed route of output port numbers from the local port to the destination.
- Using port LIDs: (Tool option `'-l'`):  
In this mode, the source and destination ports are defined by means of their LIDs. If the fabric is configured to allow multiple LIDs per port, then using any of them is valid for defining a port.
- Using port names defined in the topology file: (Tool option `'-n'`)  
This option refers to the source and destination ports by the names defined in the topology file. (Therefore, this option is relevant only if a topology file is specified to the tool.) In this mode, the tool uses the names to extract the port LIDs from the matched topology, then the tool operates as in the `'-l'` option.

## Diagnostic Utilities

The diagnostic utilities described in this chapter provide means for debugging the connectivity and status of InfiniBand (IB) devices in a fabric.

### Diagnostic Utilities

Utility	Description
<b>dumpfts</b>	Dumps tables for every switch found in an <code>ibnetdiscover</code> scan of the subnet. The dump file format is compatible with loading into OpenSM using the <code>-R file -U /path/to/dump-file</code> syntax. For further information, please refer to the tool's man page.
<b>ibaddr</b>	Can be used to show the LID and GID addresses of the specified port or the local port by default. This utility can be used as simple address resolver. For further information, please refer to the tool's man page.

Utility	Description
<b>ibcacheedit</b>	Allows users to edit an ibnetdiscover cache created through the --cache option in ibnetdiscover(8). For further information, please refer to the tool's man page.
<b>ibccconfig</b>	Supports the configuration of congestion control settings on switches and HCAs. For further information, please refer to the tool's man page.
<b>ibccquery</b>	Supports the querying of settings and other information related to congestion control. For further information, please refer to the tool's man page.
<b>ibcongest</b>	Provides static congestion analysis. It calculates routing for a given topology (topo-mode) or uses extracted lst/fdb files (lst-mode). Additionally, it analyzes congestion for a traffic schedule provided in a "schedule-file" or uses an automatically generated schedule of all-to-all-shift. To display a help message which details the tool's options, please run <code>"/opt/ibutils2/bin/ibcongest -h"</code> . For further information, please refer to the tool's man page.
<b>ibdev2netdev</b>	Enables association between IB devices and ports and the associated net device. Additionally it reports the state of the net device link. For further information, please refer to the tool's man page.
<b>ibdiagnet (of ibutils2)</b>	Scans the fabric using directed route packets and extracts all the available information regarding its connectivity and devices. An ibdiagnet run performs the following stages: <ul style="list-style-type: none"> <li>• Fabric discovery</li> <li>• Duplicated GUIDs detection</li> <li>• Links in INIT state and unresponsive links detection</li> <li>• Counters fetch</li> <li>• Error counters check</li> <li>• Routing checks</li> <li>• Link width and speed checks</li> <li>• Alias GUIDs check</li> <li>• Subnet Manager check</li> <li>• Partition keys check</li> <li>• Nodes information</li> </ul>

Utility	Description
	<p><b>Note:</b> This version of ibdiagnet is included in the ibutils2 package, and it is run by default after installing NVIDIA OFED. To use this ibdiagnet version, run: ibdiagnet. For further information, either:</p> <ol style="list-style-type: none"> <li>1. Run <code>ibdiagnet -H</code></li> </ol> <p>Or</p> <ol style="list-style-type: none"> <li>2. Refer to <a href="https://docs.nvidia.com/networking/display/ibdiagnetUserManualv10">docs.nvidia.com/networking/display/ibdiagnetUserManualv10</a></li> </ol>
<b>ibdiagnet</b>	<p>Traces a path between two end-points and provides information regarding the nodes and ports traversed along the path. It utilizes device specific health queries for the different devices along the path.</p> <p>The way ibdiagnet operates depends on the addressing mode used in the command line. If directed route addressing is used (<code>--dr_path</code> flag), the local node is the source node and the route to the destination port is known apriori (for example: <code>ibdiagnet --dr_path 0,1</code>). On the other hand, if LID-route addressing is employed, <code>--src_lid</code> and <code>--dest_lid</code>, then the source and destination ports of a route are specified by their LIDs. In this case, the actual path from the local port to the source port, and from the source port to the destination port, is defined by means of Subnet Management Linear Forwarding Table queries of the switch nodes along that path. Therefore, the path cannot be predicted as it may change.</p> <p>Example: <code>ibdiagnet --src_lid 1 --dest_lid 28</code></p> <p>For further information, please refer to the tool's -help flag.</p>
<b>ibdump</b>	<p>Dump InfiniBand traffic that flows to and from NVIDIA's ConnectX® family adapters InfiniBand ports.</p> <p>Note the following:</p> <ul style="list-style-type: none"> <li>• <code>ibdump</code> is not supported for Virtual functions (SR-IOV)</li> <li>• Infiniband traffic sniffing is supported on all HCAs</li> </ul> <p>The dump file can be loaded by the Wireshark tool for graphical traffic analysis. The following describes a workflow for local HCA (adapter) sniffing:</p> <ol style="list-style-type: none"> <li>1. Run <code>ibdump</code> with the desired options</li> <li>2. Run the application that you wish its traffic to be analyzed</li> <li>3. Stop <code>ibdump</code> (CTRL-C) or wait for the data buffer to fill (in <code>--mem-mode</code>)</li> <li>4. Open Wireshark and load the generated file</li> </ol> <p>To download Wireshark for a Linux or Windows environment go to <a href="http://www.wireshark.org">www.wireshark.org</a>.</p> <p><b>Notes:</b></p>



Utility	Description
	<ul style="list-style-type: none"> <li>• Although ibdump is a Linux application, the generated .pcap file may be analyzed on either operating system.</li> <li>• If one of the HCA's ports is configured as InfiniBand, ibdump requires IPoIB DMFS to be enabled. For further information, please refer to <a href="#">Flow Steering Configuration</a> section.</li> </ul> <p>For further information, please refer to the tool's man page.</p>
<b>iblinkinfo</b>	<p>Reports link info for each port in an InfiniBand fabric, node by node. Optionally, iblinkinfo can do partial scans and limit its output to parts of a fabric.</p> <p>For further information, please refer to the tool's man page.</p>
<b>ibnetdiscover</b>	<p>Performs InfiniBand subnet discovery and outputs a human readable topology file. GUIDs, node types, and port numbers are displayed as well as port LIDs and node descriptions. All nodes (and links) are displayed (full topology).</p> <p>This utility can also be used to list the current connected nodes. The output is printed to the standard output unless a topology file is specified.</p> <p>For further information, please refer to the tool's man page.</p>
<b>ibnetsplit</b>	<p>Automatically groups hosts and creates scripts that can be run in order to split the network into sub-networks containing one group of hosts.</p> <p>For further information, please refer to the tool's man page.</p>
<b>ibnodes</b>	<p>Uses the current InfiniBand subnet topology or an already saved topology file and extracts the InfiniBand nodes (CAs and switches).</p> <p>For further information, please refer to the tool's man page.</p>
<b>ibping</b>	<p>Uses vendor mads to validate connectivity between InfiniBand nodes. On exit, (IP) ping like output is show. ibping is run as client/server. The default is to run as client. Note also that a default ping server is implemented within the kernel.</p> <p>For further information, please refer to the tool's man page.</p>
<b>ibportstate</b>	<p>Enables querying the logical (link) and physical port states of an InfiniBand port. It also allows adjusting the link speed that is enabled on any InfiniBand port. If the queried port is a switch port, then ibportstate can be used to:</p> <ul style="list-style-type: none"> <li>• disable, enable or reset the port</li> <li>• validate the port's link width and speed against the peer port</li> </ul> <p>In case of multiple channel adapters (CAs) or multiple ports without a CA/ port being specified, a port is chosen by the utility according to the following criteria:</p>

Utility	Description
	<ul style="list-style-type: none"> <li>• The first ACTIVE port that is found.</li> <li>• If not found, the first port that is UP (physical link state is LinkUp).</li> </ul> <p>For further information, please refer to the tool's man page.</p>
<b>ibqueryerrors</b>	<p>The default behavior is to report the port error counters which exceed a threshold for each port in the fabric. The default threshold is zero (0). Error fields can also be suppressed entirely.</p> <p>In addition to reporting errors on every port, <code>ibqueryerrors</code> can report the port transmit and receive data as well as report full link information to the remote port if available.</p> <p>For further information, please refer to the tool's man page.</p>
<b>ibroute</b>	<p>Uses SMPs to display the forwarding tables—unicast (<code>LinearForwardingTable</code> or <code>LFT</code>) or multicast (<code>MulticastForwardingTable</code> or <code>MFT</code>)—for the specified switch LID and the optional lid (<code>mlid</code>) range. The default range is all valid entries in the range 1 to <code>FDBTop</code>.</p> <p>For further information, please refer to the tool's man page.</p>
<b>ibstat</b>	<p><code>ibstat</code> is a binary which displays basic information obtained from the local IB driver. Output includes LID, SMLID, port state, link width active, and port physical state.</p> <p>For further information, please refer to the tool's man page.</p>
<b>ibstats</b>	<p>Displays basic information obtained from the local InfiniBand driver. Output includes LID, SMLID, port state, port physical state, port width and port rate. For further information, please refer to the tool's man page.</p>
<b>ibswitches</b>	<p>Traces the InfiniBand subnet topology or uses an already saved topology file to extract the InfiniBand switches.</p> <p>For further information, please refer to the tool's man page.</p>
<b>ibsysstat</b>	<p>Uses vendor mads to validate connectivity between InfiniBand nodes and obtain other information about the InfiniBand node. <code>ibsysstat</code> is run as client/ server. The default is to run as client.</p> <p>For further information, please refer to the tool's man page.</p>
<b>ibtopdiff</b>	<p>Compares a topology file and a discovered listing of <code>subnet.lst/ibdiagnet.lst</code> and reports mismatches.</p> <p>Two different algorithms provided:</p>

Utility	Description
	<ul style="list-style-type: none"> <li>Using the -e option is more suitable for MANY mismatches it applies less heuristics and provide details about the match</li> <li>Providing the -s, -p and -g starts a detailed heuristics that should be used when only small number of changes are expected</li> </ul> <p>For further information, please refer to the tool's man page.</p>
<b>ibtrace</b>	<p>Uses SMPs to trace the path from a source GID/LID to a destination GID/ LID. Each hop along the path is displayed until the destination is reached or a hop does not respond. By using the -m option, multicast path tracing can be performed between source and destination nodes.</p> <p>For further information, please refer to the tool's man page.</p>
<b>ibv_asyncwatch</b>	<p>Display asynchronous events forwarded to userspace for an InfiniBand device.</p> <p>For further information, please refer to the tool's man page.</p>
<b>ibv_devices</b>	<p>Lists InfiniBand devices available for use from userspace, including node GUIDs.</p> <p>For further information, please refer to the tool's man page.</p>
<b>ibv_deviceinfo</b>	<p>Queries InfiniBand devices and prints about them information that is available for use from userspace.</p> <p>For further information, please refer to the tool's man page.</p>
<b>mstflint</b>	<p>Queries and burns a binary firmware-image file on non-volatile (Flash) memories of NVIDIA InfiniBand and Ethernet network adapters. The tool requires root privileges for Flash access.</p> <p>To run mstflint, you must know the device location on the PCI bus.</p> <p><b>Note:</b> If you purchased a standard NVIDIA network adapter card, please download the firmware image from <a href="https://www.nvidia.com/en-us/networking/support-support-firmware-download">nvidia.com/en-us/networking/Support Support Firmware Download</a>. If you purchased a non-standard card from a vendor other than NVIDIA, please contact your vendor.</p> <p>For further information, please refer to the tool's man page.</p>
<b>perquery</b>	<p>Queries InfiniBand ports' performance and error counters. Optionally, it displays aggregated counters for all ports of a node. It can also reset counters after reading them or simply reset them.</p> <p>For further information, please refer to the tool's man page.</p>

Utility	Description
<b>saquery</b>	Issues the selected SA query. Node records are queried by default. For further information, please refer to the tool's man page.
<b>sminfo</b>	Issues and dumps the output of an sminfo query in human readable format. The target SM is the one listed in the local port info or the SM specified by the optional SM LID or by the SM direct routed path. <b>Note:</b> Using sminfo for any purpose other than a simple query might result in a malfunction of the target SM. For further information, please refer to the tool's man page.
<b>smquery</b>	Sends SMP query for adaptive routing and private LFT features. For further information, please refer to the tool's man page.
<b>sm pdu mp</b>	A general purpose SMP utility which gets SM attributes from a specified SMA. The result is dumped in hex by default. For further information, please refer to the tool's man page.
<b>sm pquery</b>	Provides a basic subset of standard SMP queries to query Subnet management attributes such as node info, node description, switch info, and port info. For further information, please refer to the tool's man page.

## Link Level Retransmission (LLR) in FDR Links

With the introduction of FDR 56 Gbps technology, NVIDIA enabled a proprietary technology called LLR (Link Level Retransmission) to improve the reliability of FDR links.

This proprietary LLR technology adds additional CRC checking to the data stream and retransmits portions of packets with CRC errors at the local link level. Customers should be aware of the following facts associated with LLR technology:

- Traditional methods of checking the link health can be masked because the LLR technology automatically fixes errors. The traditional IB symbol error counter will show no errors when LLR is active.
- Latency of the fabric can be impacted slightly due to LLR retransmissions. Traditional IB performance utilities can be used to monitor any latency impact.

- Bandwidth of links can be reduced if cable performance degrades and LLR retransmissions become too numerous. Traditional IB bandwidth performance utilities can be used to monitor any bandwidth impact.

Due to these factors, an LLR retransmission rate counter has been added to the `ibdiagnet` utility that can give end users an indication of the link health.

➤ *To monitor LLR retransmission rate:*

1. Run `ibdiagnet`, no special flags required.
2. If the LLR retransmission rate limit is exceeded it will print to the screen.
3. The default limit is set to 500 and requires further investigation if exceeded.
4. The LLR retransmission rate is reflected in the results file `/var/tmp/ibdiagnet2/ibdiagnet2.pm`.

The default value of 500 retransmissions/sec has been determined by NVIDIA based on the extensive simulations and testing. Links exhibiting a lower LLR retransmission rate should not raise special concern.

## Performance Utilities

The performance utilities described in this chapter are intended to be used as a performance micro-benchmark.

Utility	Description
<b>ib_atomic_bw</b>	Calculates the BW of RDMA Atomic transactions between a pair of machines. One acts as a server and the other as a client. The client RDMA sends atomic operation to the server and calculate the BW by sampling the CPU each time it receive a successful completion. The test supports features such as Bidirectional, in which they both RDMA atomic to each other at the same time, change of MTU size, tx size, number of iteration, message size and more. Using the "-a" flag provides results for all message sizes. For further information, please refer to the tool's man page.

Utility	Description
<b>ib_atomic_latency</b>	<p>Calculates the latency of RDMA Atomic transaction of message_size between a pair of machines. One acts as a server and the other as a client. The client sends RDMA atomic operation and sample the CPU clock when it receives a successful completion, in order to calculate latency.</p> <p>For further information, please refer to the tool's man page.</p>
<b>ib_read_bw</b>	<p>Calculates the BW of RDMA read between a pair of machines. One acts as a server and the other as a client. The client RDMA reads the server memory and calculate the BW by sampling the CPU each time it receive a successful completion. The test supports features such as Bidirectional, in which they both RDMA read from each other memory's at the same time, change of MTU size, tx size, number of iteration, message size and more.</p> <p>Read is available only in RC connection mode (as specified in IB spec). For further information, please refer to the tool's man page.</p>
<b>ib_read_latency</b>	<p>Calculates the latency of RDMA read operation of message_size between a pair of machines. One acts as a server and the other as a client. They perform a ping pong benchmark on which one side RDMA reads the memory of the other side only after the other side have read his memory. Each of the sides samples the CPU clock each time they read the other side memory , in order to calculate latency. Read is available only in RC connection mode (as specified in IB spec).</p> <p>For further information, please refer to the tool's man page.</p>
<b>ib_send_bw</b>	<p>Calculates the BW of SEND between a pair of machines. One acts as a server and the other as a client. The server receive packets from the client and they both calculate the throughput of the operation. The test supports features such as Bidirectional, on which they both send and receive at the same time, change of MTU size, tx size, number of iteration, message size and more. Using the "-a" provides results for all message sizes.</p> <p>For further information, please refer to the tool's man page.</p>
<b>ib_send_latency</b>	<p>Calculates the latency of sending a packet in message_size between a pair of machines. One acts as a server and the other as a client. They perform a ping pong benchmark on which you send packet only if you receive one. Each of the sides samples the CPU each time they receive a packet in order to calculate the latency. Using the "-a" provides results for all message sizes.</p> <p>For further information, please refer to the tool's man page.</p>
<b>ib_writ</b>	<p>Calculates the BW of RDMA write between a pair of machines. One acts as a server and the other as a client. The client RDMA writes to the server memory</p>

Utility	Description
<b>e_bw</b>	and calculates the BW by sampling the CPU each time it receives a successful completion. The test supports features such as Bidirectional, in which they both RDMA write to each other at the same time, change of MTU size, tx size, number of iteration, message size and more. Using the "-a" flag provides results for all message sizes. For further information, please refer to the tool's man page.
<b>ib_write_latency</b>	Calculates the latency of RDMA write operation of message_size between a pair of machines. One acts as a server and the other as a client. They perform a ping pong benchmark on which one side RDMA writes to the other side memory only after the other side wrote on his memory. Each of the sides samples the CPU clock each time they write to the other side memory, in order to calculate latency. For further information, please refer to the tool's man page.
<b>raw_ethernet_bw</b>	Calculates the BW of SEND between a pair of machines. One acts as a server and the other as a client. The server receive packets from the client and they both calculate the throughput of the operation. The test supports features such as Bidirectional, on which they both send and receive at the same time, change of MTU size, tx size, number of iteration, message size and more. Using the "-a" provides results for all message sizes. For further information, please refer to the tool's man page.
<b>raw_ethernet_latency</b>	Calculates the latency of sending a packet in message_size between a pair of machines. One acts as a server and the other as a client. They perform a ping pong benchmark on which you send packet only if you receive one. Each of the sides samples the CPU each time they receive a packet in order to calculate the latency. Using the "-a" provides results for all message sizes. For further information, please refer to the tool's man page.

## Troubleshooting

You may be able to easily resolve the issues described in this section. If a problem persists and you are unable to resolve it yourself, please contact your NVIDIA representative or NVIDIA Support at [networking-support@nvidia.com](mailto:networking-support@nvidia.com).

The chapter contains the following sections:

- [General Issues](#)
- [Ethernet Related Issues](#)
- [InfiniBand Related Issues](#)
- [Installation Related Issues](#)
- [Performance Related Issues](#)
- [SR-IOV Related Issues](#)
- [PXE \(FlexBoot\) Related Issues](#)
- [RDMA Related Issues](#)
- [Debugging Related Issues](#)
- [OVS Offload Using ASAP2 Direct Related Issues](#)

# Common Abbreviations and Related Documents

## Common Abbreviations and Acronyms

Abbreviation/ Acronym	Description
B	(Capital) 'B' is used to indicate size in bytes or multiples of bytes (e.g., 1KB = 1024 bytes, and 1MB = 1048576 bytes)
b	(Small) 'b' is used to indicate size in bits or multiples of bits (e.g., 1Kb = 1024 bits)
FW	Firmware
HCA	Host Channel Adapter
HW	Hardware
IB	InfiniBand



Abbreviation/ Acronym	Description
iSER	iSCSI RDMA Protocol
LSB	Least significant <i>byte</i>
lsb	Least significant <i>bit</i>
MSB	Most significant <i>byte</i>
msb	Most significant <i>bit</i>
NIC	Network Interface Card
SW	Software
VPI	Virtual Protocol Interconnect
IPoIB	IP over InfiniBand
PFC	Priority Flow Control
PR	Path Record
RoCE	RDMA over Converged Ethernet
SL	Service Level
SRP	SCSI RDMA Protocol
MPI	Message Passing Interface
QoS	Quality of Service
ULP	Upper Layer Protocol
VL	Virtual Lane
vHBA	Virtual SCSI Host Bus Adapter
uDAPL	User Direct Access Programming Library

## Glossary

The following is a list of concepts and terms related to InfiniBand in general and to Subnet Managers in particular. It is included here for ease of reference, but the main reference remains the *InfiniBand Architecture Specification*.

<b>Term</b>	<b>Description</b>
Channel Adapter (CA), Host Channel Adapter (HCA)	An IB device that terminates an IB link and executes transport functions. This may be an HCA (Host CA) or a TCA (Target CA)
HCA Card	A network adapter card based on an InfiniBand channel adapter device
IB Devices	An integrated circuit implementing InfiniBand compliant communication
IB Cluster/Fabric/Subnet	A set of IB devices connected by IB cables
In-Band	A term assigned to administration activities traversing the IB connectivity only
Local Identifier (ID)	An address assigned to a port (data sink or source point) by the Subnet Manager, unique within the subnet, used for directing packets within the subnet
Local Device/Node/System	The IB Host Channel Adapter (HCA) Card installed on the machine running IBDIAG tools
Local Port	The IB port of the HCA through which IBDIAG tools connect to the IB fabric
Master Subnet Manager	The Subnet Manager that is authoritative, that has the reference configuration information for the subnet
Multicast Forwarding Tables	A table that exists in every switch providing the list of ports to forward received multicast packet. The table is organized by MLID
Network Interface Card (NIC)	A network adapter card that plugs into the PCI Express slot and provides one or more ports to an Ethernet network
Standby Subnet Manager	A Subnet Manager that is currently quiescent, and not in the role of a Master Subnet Manager, by the agency of the master SM
Subnet Administrator (SA)	An application (normally part of the Subnet Manager) that implements the interface for querying and manipulating subnet management data
Subnet Manager (SM)	One of several entities involved in the configuration and control of the IB fabric

Term	Description
Unicast Linear Forwarding Tables (LFT)	A table that exists in every switch providing the port through which packets should be sent to each LID
Virtual Protocol Interconnect (VPI)	An NVIDIA technology that allows NVIDIA channel adapter devices (ConnectX®) to simultaneously connect to an InfiniBand subnet and a 10GigE subnet (each subnet connects to one of the adapter ports)

## Related Documentation

Document Name	Description
InfiniBand Architecture Specification, Vol. 1, Release 1.2.1	The InfiniBand Architecture Specification that is provided by IBTA
IEEE Std 802.3ae™-2002 (Amendment to IEEE Std 802.3-2002) Document # PDF: SS94996	Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications Amendment: Media Access Control (MAC) Parameters, Physical Layers, and Management Parameters for 10 Gb/s Operation
Firmware Release Notes for NVIDIA adapter devices	See the <a href="#">Release Notes</a> relevant to your adapter device
MFT User Manual and Release Notes	NVIDIA Firmware Tools (MFT) User Manual and Release Notes documents
WinOF User Manual	Mellanox WinOF User Manual describes the installation, configuration, and operation of NVIDIA Windows driver
VMA User Manual	NVIDIA VMA User Manual describes the installation, configuration, and operation of NVIDIA VMA driver

---

# Documentation History

- [Release Notes History](#)
- [User Manual Revision History](#)

## Release Notes History

- [Changes and New Features History](#)
- [Bug Fixes History](#)

## User Manual Revision History

Release	Date	Description
24.04	April 2024	Added <a href="#">SR-IOV Live Migration</a>
23.10	November 2023	<ul style="list-style-type: none"><li>• Added <a href="#">IPsec Full Offload</a></li><li>• Added a note to <a href="#">NVME-oF - NVM Express over Fabrics</a>.</li></ul>
23.07	September 2023	No changes to the User Manual.
	August 2023	<ul style="list-style-type: none"><li>• Updated "Interrupt Request (IRQ) Naming" under <a href="#">Ethernet Interface</a>.</li><li>• Removed the "IPsec Crypto Offload" section from the <a href="#">Ethernet Network</a> chapter.</li><li>• Removed all "Connected Mode" indications, and the "IPoIB Mode Setting" section from <a href="#">IP over InfiniBand (IPoIB)</a>.</li><li>• Removed Running OpenSM, and updated Running OpenSM as Daemon under <a href="#">NVIDIA SM</a>.</li></ul>

Release	Date	Description
23.04	May 2023	<ul style="list-style-type: none"> <li>• Removed "Installing NEO-Host Using mlnxofedinstall Script" subsection from the <a href="#">Installing MLNX_OFED</a> section</li> <li>• Added note under "<a href="#">Enrolling NVIDIA's x.509 Public Key On your Systems</a>" subsection in the <a href="#">UEFI Secure Boot</a> section</li> <li>• Updated the <a href="#">NVIDIA SM (OpenSM)</a> section</li> </ul>
	Jun 2023	<ul style="list-style-type: none"> <li>• Added an important note to <a href="#">IPsec Crypto Offload</a></li> </ul>
5.7	August 2022	<ul style="list-style-type: none"> <li>• Added <a href="#">Out of Order (OOO)</a> under RoCE section</li> <li>• Added section describing <a href="#">Dumping Steering Info</a></li> </ul>
5.6	April 2022	<ul style="list-style-type: none"> <li>• Added <a href="#">OpenFlow Meters</a> section</li> <li>• Added <a href="#">Installation on Community Operating Systems</a> section</li> </ul>

---

# Legal Notices and 3rd Party Licenses

The following are the drivers' software, tools and HCA firmware legal notices and 3rd party licenses.

Product	Version	Legal Notices and 3rd Party Licenses
MLNX_OFED	24.04-0.6.6.0	<ul style="list-style-type: none"><li>• <a href="#">License</a></li><li>• <a href="#">3rd Part Notice</a></li></ul>
Firmware	xx.41.10xx	<ul style="list-style-type: none"><li>• <a href="#">HCA Firmware EULA</a></li><li>• <a href="#">3rd Party Unify Notice</a></li><li>• <a href="#">License</a></li></ul>
MFT	4.28.0	<ul style="list-style-type: none"><li>• <a href="#">3rd Party Notice</a></li><li>• <a href="#">License</a></li></ul>
<b>Clusterkit</b>	1.13.461	<ul style="list-style-type: none"><li>• <a href="#">License</a></li><li>• <a href="#">3rd Party Unify Notice</a></li></ul>
DPCP	1.1.48	<ul style="list-style-type: none"><li>• <a href="#">License</a></li><li>• <a href="#">3rd Party Notice</a></li></ul>
VMA	9.8.60	<ul style="list-style-type: none"><li>• <a href="#">License</a></li><li>• <a href="#">3rd Party Notice</a></li></ul>
XLIO	3.30.5	<ul style="list-style-type: none"><li>• <a href="#">License</a></li><li>• <a href="#">3rd Party Notice</a></li></ul>
HCOLL	4.8.3227	<ul style="list-style-type: none"><li>• <a href="#">License</a></li><li>• <a href="#">3rd Party Unify Notice</a></li></ul>

Product	Version	Legal Notices and 3rd Party Licenses
SHARP	3.7.0	<ul style="list-style-type: none"> <li>• <a href="#">License</a></li> <li>• <a href="#">3rd Party Unify Notice</a></li> </ul>
ibutils2	2.17	<ul style="list-style-type: none"> <li>• <a href="#">3rd Party Notice</a></li> <li>• <a href="#">3rd Party Unify Notice</a></li> </ul>
OpenSM	5.19.0	<ul style="list-style-type: none"> <li>• <a href="#">3rd Party Unify Notice</a></li> <li>• <a href="#">3rd Party Notice</a></li> </ul>
mpitests	3.2.23	<ul style="list-style-type: none"> <li>• <a href="#">License</a></li> <li>• <a href="#">3rd Party Unify Notice</a></li> </ul>

© Copyright 2024, NVIDIA. PDF Generated on 06/05/2024